CHAPTER

3

THE SYSTEM BOARD

In this chapter, you will learn:

- Which physical components are on the system board
- How the system board transports data, follows programming logic, and coordinates the timing and execution of each processing task
- About the recent evolution of several system-board components
- How to set the CPU and system bus frequency for the system board

Chapters 1 and 2 surveyed the hardware and software that make up a personal computer and described how they work together to create a functioning computer system. In this chapter, we begin to examine in detail how the components of a computer work in harmony and with accuracy. Our starting point is the system board, the central site of computer logic circuitry and the location of the most important microchip in the computer, the CPU.

To understand the ideas in this chapter, you should (1) know the definitions of bit, byte, kilobyte, and hexadecimal (hex) number, and (2) be able to read memory addresses written in hex. Appendix D describes bits, bytes, and hex numbers used to address memory locations. If this is unfamiliar territory for you, turn to Appendix D before reading on.

Types of System Boards

A+CORE 1.1, 4.3

A system board's primary purpose is to house the CPU and allow all devices to communicate with it and each other. The two most popular system boards are the older AT and the newer ATX. The AT system board has a power connection for 5- and 12-volt lines coming from the power supply. To accommodate the newer CPUs that use less voltage, the ATX has lines for 5, 12, and 3.3 volts from the power supply. Figure 3-1 shows that the ATX system board uses a single P1 power connection, but the AT board uses two power connections, P8 and P9.

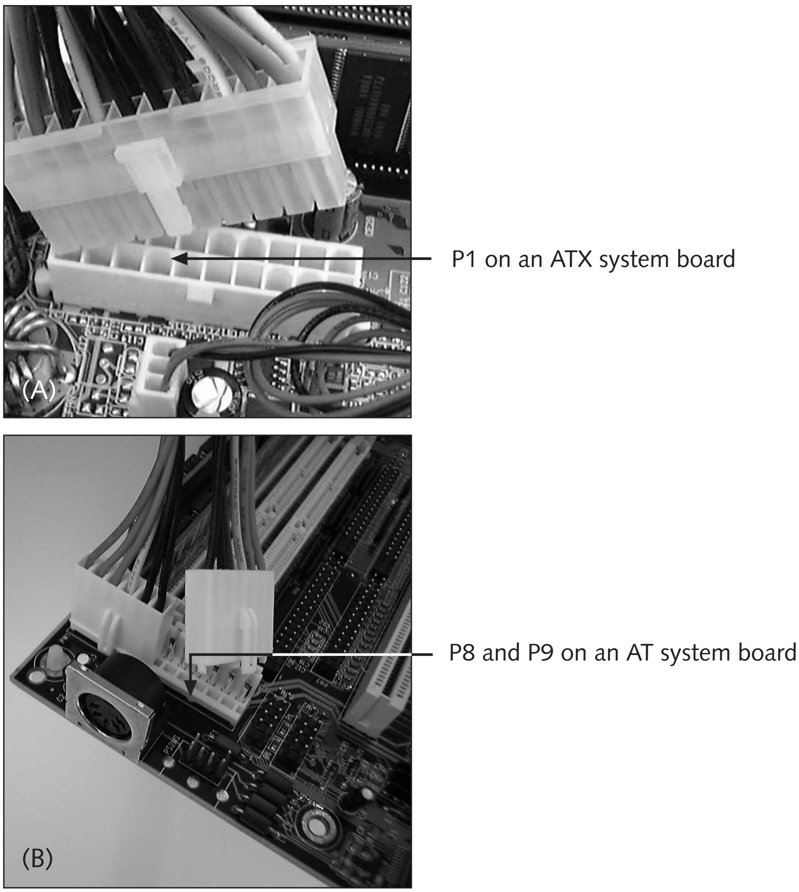

Figure 3-1 ATX uses a single P1 power connection (A), but AT type system boards use P8 and P9 power connections

Each board is available in two sizes. The ATX boards include more power-management features and support faster systems. Table 3-1 summarizes these different boards and their form factors. ("Form factor" is computer jargon for the size and shape of a board or other device.)

A+CORE 1.1, 4.3

Table 3-1 Types of system boards

Types of System Boards	Description
AT	■ Oldest type of system board still commonly used ■ Uses P8 and P9 power connections (See Figure 3-1) ■ Measures 30.5 cm × 33 cm
Baby AT	■ Smaller version of AT. Small size is possible because system-board logic is stored on a smaller chip set. ■ Uses P8 and P9 power connections ■ Measures 33 cm × 22 cm
ATX	■ Developed by Intel for Pentium systems ■ Has a more conveniently accessible layout than AT boards ■ Includes a power-on switch that can be software-enabled and extra power connections for extra fans ■ Uses a single 20-pin power connection called a **P1 connector** (See Figure 3-1) ■ Measures 30.5 cm × 24.4 cm
Mini ATX	■ An ATX board with a more compact design ■ Measures 28.4 cm × 20.8 cm

The main components on a system board are the following:

- CPU and its accompanying chip set
- System clock
- ROM BIOS
- CMOS configuration chip and its battery
- RAM
- RAM cache (only on older system boards)
- System bus with expansion slots
- Jumpers and DIP switches
- Ports that come directly off the board
- Power supply connections

A+CORE 4.3

Of the components listed above, you can replace or upgrade the following five: CPU, ROM BIOS chip, CMOS battery, RAM, and RAM cache. Because you can exchange these items without returning the system board to the manufacturer, they are called **field replaceable units**.

Before examining the most important system board components, let's look at the system board itself as a component (see Figures 3-2 and 3-3).

A+ CORE 1.1

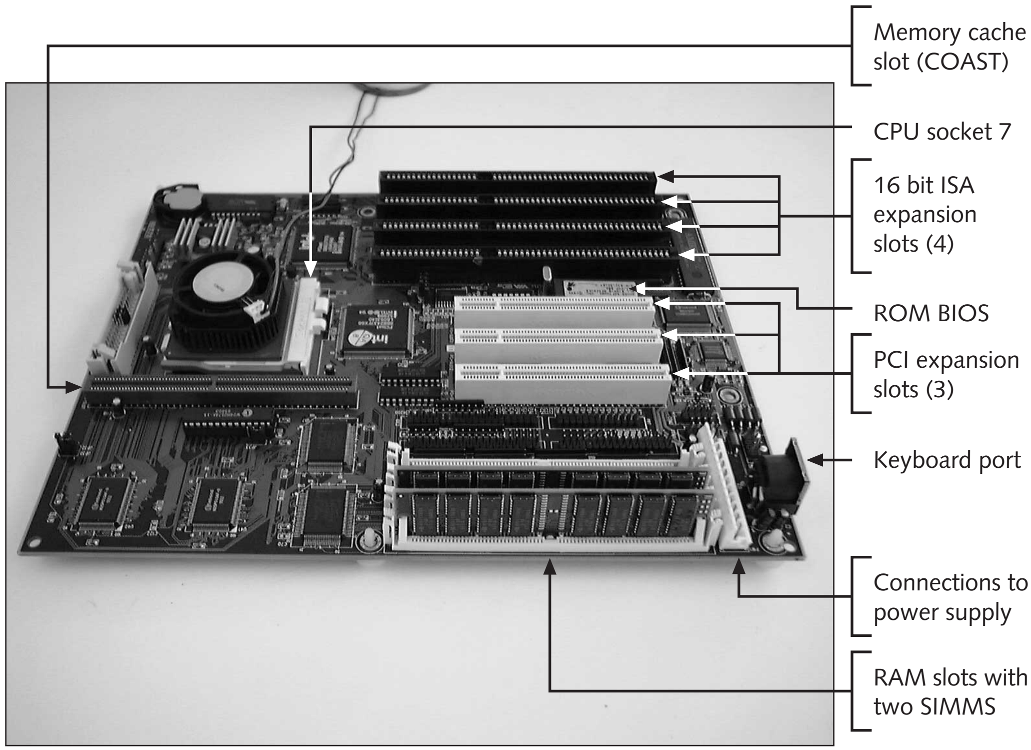

Figure 3-2 A typical AT system board with memory cache and socket 7 for the Intel Classic Pentium CPU. The CPU with a fan on top is installed as well as two SIMM memory modules

When you buy a system board, your selection determines the following components:

- Types and speeds of the CPU you can use
- Chip set on the board (already installed)
- Memory cache type and size
- Types and number of expansion slots: EISA, PCI, and AGP (explained below)
- Type of memory: ECC, EDO, SDRAM, SIMMs, or DIMMs (explained below)
- Maximum amount of memory you can install on the board and the incremental amounts by which you can upgrade memory
- Type of case you can use
- ROM BIOS (already installed)
- Type of keyboard connector
- Presence or absence of different types of proprietary video and/or proprietary local bus slots

- Presence or absence of IDE adapters and SCSI controller (explained below)
- Presence or absence of COM ports, LPT ports, and mouse port

3

A+CORE 1.1

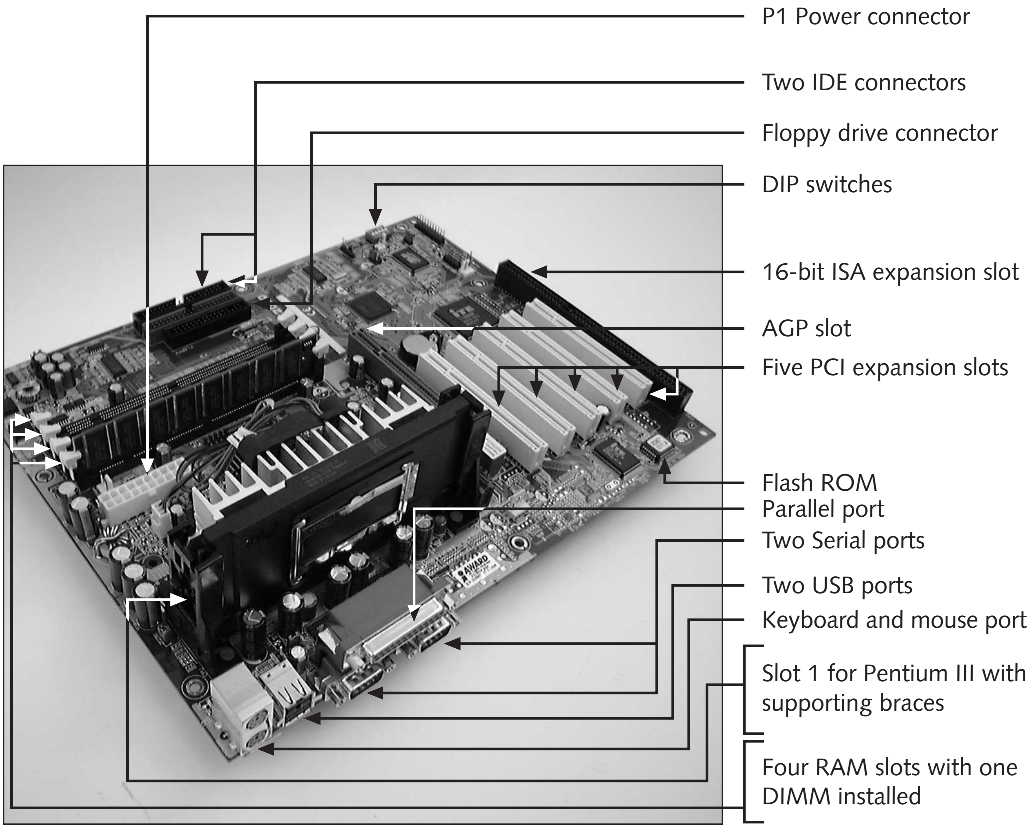

Figure 3-3 An ATX system board with a Pentium III and one DIMM module installed

Selecting the system board is, therefore, a very important decision when you purchase a computer or assemble one from parts, because the system board determines so many of your computer's features.

Depending on which applications and peripheral devices you plan to use with the computer, you can take one of three different approaches to selecting a system board. The first option is to select the board that provides the most room for expansion so you can upgrade and exchange components and add-on devices easily. A second approach is to select the board that best suits the needs of the computer's current configuration, knowing that when you need to upgrade, you will likely switch to new technology and a new system board. The third approach is to select a system board that meets your present needs with moderate room for expansion.

Ask the following questions when selecting a system board:

- Is the system board designed so that long expansion cards don't get in the way of the CPU or other important devices you might want to access?

- How many different CPUs can the system board support—only those manufactured by Intel or also those made by Intel's competitors?
- What bus speeds, type of memory, and system BIOS does the board support?
- Does the board use many embedded devices (discussed below)?
- Does the board fit the case I plan to use?
- Does the board support my legacy cards?
- What is the warranty on the board?
- How extensive and user-friendly is the documentation?
- How much support does the manufacturer supply for the board?

Sometimes a system board contains a component that is more commonly offered as a separate device. A component on the board is called an embedded component. One example is support for video. The video port might be on the system board or might require a video card. The cost of a system board with an embedded component is usually less than the combined cost of a system board with an expansion card but no component. If you plan to expand, be cautious about choosing a proprietary board that has many embedded components. A proprietary design using many embedded devices often does not easily accept add-on devices from other manufacturers. For example, if you plan to add a more powerful video card, you might not want to choose a system board that contains an embedded video controller.

Even though you can often set a switch on the system board to disable the proprietary video controller, there is little advantage to paying the extra money for the on-board video controller.

If you have an embedded component, make sure you can disable the component so that you can use another external component if needed. You disable a component on the system board through jumpers on the board or through CMOS setup.

Table 3-2 lists some manufacturers of system boards with their web addresses.

Table 3-2 Major manufacturers of system boards

Manufacturer	Web Address
motherboard.com	www.motherboards.com
American Megatrends, Inc.	www.megatrends.com
ASUS	www.asus.com
Diamond Multimedia	www.diamondmm.com
First International Computer, Inc.	www.fica.com
Giga-Byte Technology Co., Ltd.	www.giga-byte.com
Intel Corporation	www.intel.com
Supermicro Computer, Inc.	www.supermicro.com
Tyan Computer Corporation	www.tyan.com

THE SYSTEM CLOCK

Remember from Chapter 1 that the system board contains a system clock that keeps the beat for many system-board activities. We use units called megahertz (MHz) to measure clock frequency. One megahertz (MHz) is equal to 1,000,000 beats, or cycles, of the clock per second. A single clock beat or cycle was once the smallest unit of processing the CPU or another device could execute, meaning that it could only do one thing for each beat of the clock. Some CPUs today can perform two activities per clock cycle. Even though how fast a CPU can operate is often referred to as the CPU speed, it is more accurate but less common to speak of the CPU frequency. For example, you might say that a CPU can operate at a frequency of 550 MHz.

A wait state occurs when the CPU must wait for another component, for example when slower dynamic RAM reads or writes data. To allow time for the slow operation, CMOS setup information specifies that the CPU maintain a wait state. If the CPU normally can do something in two clock beats, for example, it is told to wait an extra clock beat, meaning its cycle takes a total of three clock beats. It works for two beats and then waits one beat, which makes for a 50% slowdown. Wait states might be incorporated to slow the CPU so that the rest of the system-board activity can keep up. Wait states are initially set as part of the system board's default settings and are only changed in rare circumstances, such as when the board becomes unstable.

THE CPU AND THE CHIP SET

IBM and IBM-compatible computers manufactured today use a microprocessor chip made by Intel or one of its competitors. Early CPUs by Intel were identified by model numbers: 8088, 8086, 80286, 386, and 486. The next CPU introduced after the 486 was named the Pentium, and all Intel CPUs after that include Pentium in their name. The model numbers can be written with or without the 80 prefix and are sometimes preceded with an i as in 80486, 486, or i486. The name Pentium comes from the word *pente*, the Greek word for five, which Intel chose after a legal battle with two competitors, AMD and Cyrix. AMD and Cyrix won rights to continue using the X86 chip names, but are not allowed to use the word "Pentium" to name their CPUs.

You need to know how to identify a CPU installed in a system and what performance to expect from that CPU. The following attributes are used to rate CPUs:

A+CORE 1.1, 4.1

1. **CPU speed measured in megahertz**. The first CPU used in an IBM PC was the 8088, which worked at about 4.77 MHz, or 4,770,000 clock beats per second. An average speed for a new CPU today is about 550 MHz, or 550,000,000 beats per second. In less than a minute this processor beats more times than your heart beats in a lifetime!
2. **Efficiency of the programming code**. Permanently built into the CPU chip are programs that accomplish fundamental operations, such as how to compare or add two numbers. Less efficient CPUs require more steps to perform these simple operations than more efficient CPUs. These groups of instructions are collectively called the "instruction set."

A+ CORE 1.1, 4.1

3. **Word size, sometimes called the internal data path size**. Word size is the largest number of bits the CPU can process in one operation. Word size ranges from 16 bits (2 bytes) to 64 bits (8 bytes).
4. **Data path**. The data path, sometimes called the external data path size, is the largest number of bits that can be transported into the CPU. The size of the data path is the same as the system bus size, or the number of bits that can be transported along the bus at one time. (The data path ranges from 8 bits to 64 bits.) The word size need not be as large as the data path size; some CPUs can receive more bits than they can process at one time.
5. **Maximum number of memory addresses**. A computer case has room for a lot of memory physically housed within the case, but a CPU has only a fixed range of addresses that it can assign to this physical memory. How many memory addresses the CPU can assign limits the amount of physical memory chips that the computer can effectively use. The minimum number of memory addresses a CPU can use is one megabyte (where each byte of memory is assigned a single address). Recall that one megabyte is equal to 1024 kilobytes, which is equal to 1024 $\times$ 1024 bytes, or 1,048,576 memory addresses. The maximum number of memory addresses for Pentium CPUs is 4096 megabytes, which is equal to 4 gigabytes.
6. **The amount of memory included with the CPU**. Some CPUs have storage for instructions and data built inside the chip housing. This is called internal cache, primary cache, level 1, or L1 cache.
7. **Multiprocessing ability**. Some microchips are really two processors in one and can do more than one thing at a time. Others are designed to work in cooperation with other CPUs installed on the same system board.
8. **Special functionality**. Special purpose CPUs, such as the Pentium MMX CPU, which is designed to manage multimedia devices efficiently.

Until Intel manufactured the Pentium series of chips, the three most popular ways of measuring CPU power were speed measured in megahertz and word size and data path size measured in bits. Our criteria for measuring the power of a CPU have changed since the introduction of the Pentium. The word size and path size are no longer distinguishing qualities, because these sizes have not changed significantly for the last few years. Currently, we are more interested in clock speed, bus speed, internal cache, and especially, the intended functionality of the chip, such as its ability to handle graphics well (MMX technology).

Relating CPU Attributes to Bus Architecture

A+ CORE 4.1, 4.3

Two of the CPU attributes listed above work in relation to the bus architecture: number of memory addresses and data path size. The data path size is determined by the width of the bus data path, or the number of parallel wires in the bus data path, and the number of memory addresses is determined by the number of traces, or wires, on the bus that are used for memory addresses.

A+CORE 4.1, 4.3

Recall that if a data path is 16-bits wide, there are 16 wires on the bus, each used to transmit one bit, and 16 pins connecting to the CPU that can input and output single bits. If a bus has 20 wires dedicated to memory addresses transmitted over the bus, the CPU can transmit a maximum of 20 bits to define one memory address. The largest 20-bit base 2 number possible is then the maximum number of memory addresses the CPU can use. That number is 1111 1111 1111 1111 1111 in binary, or 1 MB of memory addresses (1,048,576 unique addresses).

3

The Earlier Intel CPUs

Table 3-3 lists specifications for some early CPUs made by Intel. Until the introduction of Pentium chips and their clones, most chips were rated by the criteria listed in this table.

Table 3-3 The power of the early Intel CPUs

Model (chronological order)	Approximate Speed (MHz)	Word Size (bits)	Path Size (bits)	Memory Addresses (MB)
80386DX	40	32	32	4096
80386SX	33	32	16	16
486DX	60	32	32	4096
486SX	25	32	32	4096
First Pentium	60	32×2	64	4096

Looking at the first two rows of Table 3-3, note that the 80386SX chip had a smaller path size than the 80386DX, although it was developed later. At the time Intel first manufactured the 80386DX with its 32-bit path size, system-board manufacturers could produce at a reasonable cost a system board with a path size of only 16 bits, or 2 bytes. Therefore the system-board manufacturers could not take advantage of the DX's 32-bit path size and chose not to use the first 80386DX chips. In response to this, Intel produced the cheaper 80386SX chip, which accommodated the smaller path size and kept the cost of the system more reasonable for personal computer users. The 80386SX chip used an internal 32-bit word size but an external 16-bit path size. (Internal refers to operations inside the CPU, and external refers to operations between the inside and outside of the CPU, such as those on the bus.) The smaller path size of the 80386SX is the reason that it is slower than the 80386DX chip (S stands for single and D for double).

Table 3-3 lists the earlier CPUs chronologically, based on their introduction in the marketplace. If you look at one of these CPUs, you see it is labeled as 80386SX-16, 80486DX2-50, or another number using a similar convention. The number at the end of the model number, 16 or 50 in the examples, refers to the speed of the CPU in megahertz. The 2 following the 486DX CPU indicates that the chip can work in overdrive mode, which doubles the external clock speed to increase the overall speed of the computer. (On some older computers, doubling the clock speed was called **turbo mode** and was accomplished by pressing a button on

the computer case.) Sometimes, system boards and CPUs that work in this overdrive mode overheat, and heat sinks and/or fans must be mounted on top of the CPU.

For notebook computers, the CPU model number often has an L in it, as in 486SL-20. The L indicates that this microchip is a 486SX that requires a lower voltage than the regular SX. The 20 indicates that the speed is 20 megahertz.

Voltages Used by CPUs

Early CPUs all used 5 volts of electrical current to operate; these included the 486DX, 80486SX, 80487SX, and 80486DX2. Later versions of the 80486SX and the 80486DX4 CPUs ran on 3.3 volts. Because the power supply to the system board only supplied 5 and 12 volts, a voltage regulator was used to provide the 3.3 volts. The first Pentium running at 60/66 MHz used 5 volts. All other Pentiums, including the Pentium Pro and Pentium II, use 3.3 volts and 2.8 volts.

Coprocessor Used with Older CPUs

Some older CPU microchips were designed to work hand in hand with a secondary microchip processor called a coprocessor. The coprocessor performed calculations for the CPU at a faster speed than the CPU could. The coprocessor for the 80386 chip is the 80387. The 486DX has the coprocessor built into the CPU housing. The 486SX has the coprocessor portion of the chip disabled. Software must be written to make use of a coprocessor. Most software today assumes that you have a 486DX or Pentium chip and writes its code to take advantage of this coprocessor capability.

The Pentium and Its Competitors

A+ CORE 4.1

The latest CPU microchips by Intel are the Pentium series of chips. A Pentium chip has two arithmetic logic units, meaning that it can perform two calculations at the same time; it is therefore a true multiprocessor. Pentiums have a 64-bit external path size and two 32-bit internal paths, one for each arithmetic logic unit.

Comparing Chips

To compare the Pentium family of chips and the Pentium competitors, you need to understand bus speed, processor speed, multiplier, and memory cache. Each of these is introduced here and then discussed in more detail later in the chapter.

A+ CORE 4.1

Recall that **bus speed** is the frequency or speed at which data moves on a **bus**. Remember also that a system board has several buses; later in the chapter you will learn the details of each. Each bus runs at a certain speed, some faster than others. Only the fastest bus connects directly to the CPU. This bus goes by many names. It's called the system-board bus, or the system bus, because it's the main bus on the system board connecting directly to the CPU, or the Pentium bus because it connects directly to the Pentium. It's called the host bus because other buses connect to it to get to the CPU, and it's also called the **memory bus**

A+CORE 4.1

because it connects the CPU to RAM. This book uses the name memory bus because it is the most descriptive of the four, although system bus is probably the more popular term. However, the term system bus is used so loosely in the literature that its meaning can be confusing. The common speeds for the memory bus are 66 MHz, 75 MHz, 100 MHz, 133 MHz, and 200 MHz, although the bus can operate at several other speeds, depending on how jumpers are set on the system board.

TIP

When you read that Intel supports a system-board speed of 66 and 100 MHz, and that its competitors support bus speeds of 75 MHz, these speeds are all referencing the memory bus speed. In documentation you sometimes see the memory bus speed called the bus clock because the pulses generated on the clock line of the bus determine its speed. Other slower buses connect to the memory bus, which serves as the go-between for other buses and the CPU.

Processor speed is the speed at which the CPU is operating internally. If the CPU operates at 150 MHz internally, but 75 MHz externally, the processor speed is 150 MHz and the memory bus speed is 75 MHz. The CPU is operating at twice the speed of the bus. This factor is called the **multiplier**. If you multiply the memory bus speed by the multiplier, you get the processor speed or the speed of the CPU:

Memory bus speed × multiplier = processor speed

You can use jumpers on the system board to set the memory bus speed or bus clock. The jumpers set the multiplier, which then determines the CPU speed or processor speed. Common multipliers are 1.5, 2, 2.5, 3, 3.5, and 4.

A+CORE 4.1, 4.3

A **memory cache** is a small amount of RAM (referred to as static RAM or SRAM) that is much faster than the rest of RAM, which is called dynamic RAM (DRAM) because it loses its data rapidly and must be refreshed often. Refreshing RAM takes time, making DRAM slower than SRAM, which does not need refreshing because it can hold its data as long as power is available. Therefore, both programming code and data can be stored temporarily in this faster static RAM cache to speed up the CPU processing of both. The size of the cache a CPU can support is a measure of its performance, especially during intense calculations.

A memory cache that is included on the CPU microchip itself is called **internal cache**, **primary cache**, **Level 1**, or L1 cache. A cache outside of the CPU microchip is called **external cache**, secondary cache, **Level 2**, or L2 cache. L2 caches are usually 128K, 256K, 512K, or 1 MB in size. In the past, all L2 cache was contained on the system board, but beginning with the Pentium Pro, some L2 cache has been included inside the Pentium physical housing—not on the CPU microchip like the L1 cache, but on a tiny circuit board with the CPU chip, within the same housing. The bus between the processor and the L2 cache is called the **backside bus** or cache bus and is not visible, because it is completely contained inside the CPU housing (see Figure 3-4). On the Pentium Pro and Pentium II, this cache bus runs at half the speed of the processor.

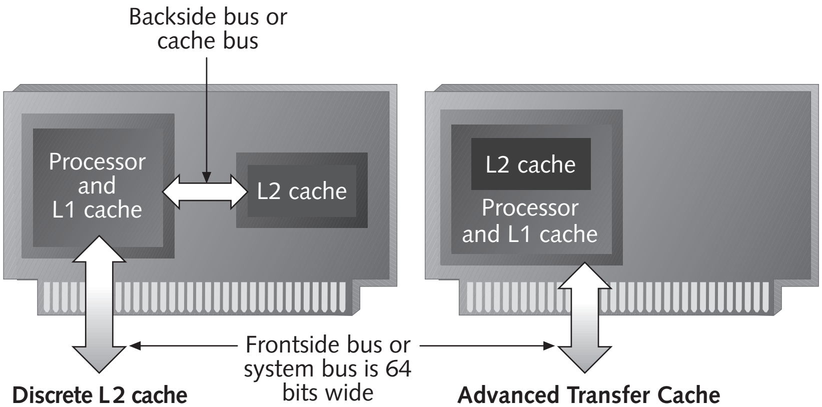

Figure 3-4 Some Pentiums contain L2 cache on separate dies (discrete L2 cache), and some contain L2 cache on the same die (Advanced Transfer cache)

A+CORE 4.1

In contrast, the bus that connects the CPU to memory outside the housing is called the **frontside bus** and can be seen on the system board. The frontside bus is the same bus as the memory bus. (Now the memory bus has one more name!)

More recently, some Pentium III CPUs contain L2 cache directly on the same die as the processor core and this is called **Advanced Transfer Cache (ATC)**, making it difficult to distinguish between L1 and L2 cache. ATC makes it possible for the Pentium III to fit on a smaller and less expensive form factor (the shape and size of a device). The ATC Cache bus is 256 bits wide and runs at the same speed as the processor. Pentium III L2 cache stored on a separate microchip within the CPU housing is called **discrete L2 cache** (see Figure 3-4). With discrete L2 cache, the Pentium III cache bus is 64 bits wide and runs at half the speed as the processor. All Pentium III processors have either 512K of discrete cache or 256K of ATC cache.

Table 3-4 lists the five types of Pentium CPUs, the Classic Pentium, Pentium MMX, Pentium Pro, Pentium II, and Pentium III. Variations of the Pentium II processor include the Celeron and Xeon. Each one is discussed below.

Table 3-4 The Intel Pentium family of CPUs

Processor	Current Processor Speeds (MHz)	Primary L1 Cache	Secondary L2 Cache	System Bus Speeds (MHz)
Classic Pentium	60, 66, 75, 90, 100, 120, 133, 150, 166, 200	16K	None	66
Pentium MMX	133, 150, 166, 200, 233, 266	32K	None	66
Pentium Pro	150, 166, 180, 200	16K	256K, 512K, or 1 MB	60, 66
Pentium II	233, 266, 300, 333, 350, 366, 400, 450	32K	256K, 512K	66, 100

A+CORE 4.1

Table 3-4 The Intel Pentium family of CPUs (continued)

Processor	Current Processor Speeds (MHz)	Primary L1 Cache	Secondary L2 Cache	System Bus Speeds (MHz)
Celeron	266, 300, 333, 366, 400, 433, 450 (mobile) 466, 500, 533, 566, 600	32K	Some have 128K	66, 100 (mobile only)
Pentium II Xeon	400, 450	32K	512K, 1 MB, or 2 MB	100
Pentium III	400, 450, 500, 533, 550, 600, 650, 667, 700, 733, 750, 800, 850, 866, 933, 1 GHz	32K	256, 512	100, 133
Pentium III Xeon	500, 550, 600, 667, 700, 733, 800, 866	32K	256K to 2 MB	100, 133

3

Classic Pentium

The first Pentium chip was introduced in March 1993, and has now become affectionately known as the "Classic Pentium." Early problems with this first Pentium (which Intel later resolved) could cause errors such as incorrect calculations on spreadsheets. The Classic Pentium is no longer manufactured.

Pentium MMX

The Pentium MMX (Multimedia Extension) targets the home market. It speeds up graphical applications and performs well with games and multimedia software.

Pentium Pro

Intel recommends the Pentium Pro for 32-bit applications that rely heavily on fast access to large amounts of cache memory. It was the first Pentium to offer Level 2 cache inside the CPU housing as well as other features not available on the Classic Pentium. The Pentium Pro is popular for computing-intensive workstations and servers but, because it does not perform well in real mode, it does not perform well with older 16-bit application software written for DOS or Windows 3.x.

Pentium II

The Pentium II is designed for graphics-intensive workstations and servers, and works well with 3-D graphic manipulation, CAD (computer-aided design), and multimedia presentations. The Pentium II is the first Pentium to use a slot (Slot 1) instead of a socket

A+CORE 4.1

to connect to the system board. (CPU sockets and slots are covered later in the chapter.) Intel chose to patent Slot 1, and, in doing so, forced its competitors to stay with the slower socket technology as they developed equivalent processors. The Pentium II can use the 100-MHz memory bus with processor speeds up to 450 MHz.

The Celeron processor is a low-end Pentium II processor that targets the low-end multimedia PC market segment. It uses Level 2 cache within the processor housing and works well with Windows 9x and the most common applications.

The Pentium II Xeon processor is a fast, high-end Pentium II processor designed exclusively for servers and powerful workstations. It can support up to eight processors in one computer and is recommended for use with Windows NT, Windows 2000, and UNIX operating systems.

Pentium III

The Pentium III (see Figure 3-5) uses either a slot or a socket and runs with the 100-MHz or 133-MHz memory bus with a processor speed up to 1 GHz. The Pentium III introduced Intel's new performance enhancement called SSE, for Streaming SIMD Extensions. (SIMD stands for single instruction, multiple data, and is a method used by MMX to speed up multimedia processing.) SSE is a new instruction set designed to improve multimedia processing even further. SSE will be an improvement over MMX as soon as operating systems and applications software are written to use it.

Figure 3-5 This Pentium III is contained in a SECC cartridge that stands on its end in Slot 1 on a system board

A+CORE 4.1

The Pentium III Xeon is a high-end Pentium III processor that runs on the 133 MHz system bus and is designed for mid range servers and high-end workstations. It uses a 330-pin slot called the SC330 (slot connector 330), sometimes called Slot 2, and is contained within a cartridge called a Single Edge Contact Cartridge (SECC).

The Pentium Competitors

Intel's two primary competitors are AMD and Cyrix. Both companies have advertised goals to produce CPUs that are just as fast and powerful as Intel's, but at a lower cost. For the latest information about the Pentium and its competitors, see these web sites: *www.amd.com*, *www.cyrix.com*, and *www.intel.com*.

Table 3-5 lists the two early processors that competed with Intel's Classic Pentium. Neither of these processors is manufactured today, although plenty of them are still in use.

Table 3-5 Cyrix and AMD competitors of the Classic Pentium

Processor	Current Processor Speeds (MHz)	Bus Speeds (MHz)	Multiplier	Internal or Primary Cache
Cyrix 6x86 or M1	150	75	2	16K
AMD K5	75, 90, 100, 116, 133	50, 60, 66	1.5 or 1.75	24K

The AMD K5 offers an unusual assortment of clock speeds and bus speeds. One disadvantage of the Cyrix 6x86 is that it uses an external bus speed of 75 MHz, which is not supported by Intel for its chip set (that is, Intel does not guarantee this bus speed to be stable). Therefore, if a system uses this Cyrix chip, the system board must use another brand of chip set other than the popular Intel brand or, if the Intel chip set is used, the system board must be set to run at a bus speed that is not guaranteed by Intel to be stable.

Running a system board at a higher speed than that suggested by the manufacturer is called **overclocking** and is not recommended, because the speed is not guaranteed to be stable by Intel. VIA and SiS both have chip sets that support the memory bus speeds needed by the Cyrix CPUs.

Competitors of the Advanced Pentiums

Table 3-6 shows the performance ratings of five competitors of the Pentium advanced processors. Cyrix processors, such as the Cyrix III shown in Figure 3-6, use sockets that can also be used by Intel Pentium processors, but AMD has taken a different approach. AMD processors that can run on a 100-MHz system bus use a special type of socket called Super Socket 7 that supports an AGP video slot and 100 MHz system bus. AGP refers to a special port for video cards called the accelerated graphics port (AGP), which is discussed later in the chapter. The AMD Athlon, shown in Figure 3-7, uses a proprietary 242-pin slot called Slot A that looks like the Intel Slot 1, which also has 242 pins.

A+CORE 4.1

Table 3-6 Cyrix and AMD competitors of the advanced Pentiums

Processor	Current Clock Speeds (MHz)	Compares to	System Bus Speed (MHz)	Socket or Slot
Cyrix M II	300, 333, 350	Pentium II, Celeron	66, 75, 83, 95, 100	Socket 7
Cyrix III	433, 466, 500, 533	Celeron, Pentium III	66, 100, 133	Socket 370
AMD-K6-2	166, 200, 266, 300, 333, 350, 366, 380, 400, 450, 475	Pentium II, Celeron	66, 95, 100	Socket 7 or Super Socket 7
AMD-K6-III	350, 366, 380, 400, 433, 450	Pentium II	100	Super Socket 7
AMD Athlon	600, 650, 700, 750, 800, 850, 900, 950, 1 GHz	Pentium III	200	Slot A

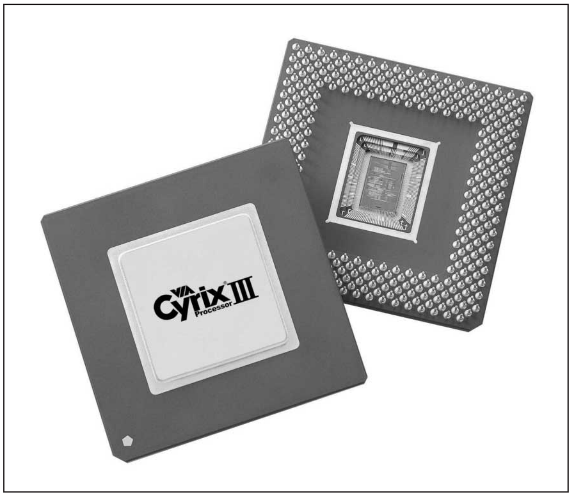

Figure 3-6 The Cyrix III competes with the Intel Celeron and the Pentium III

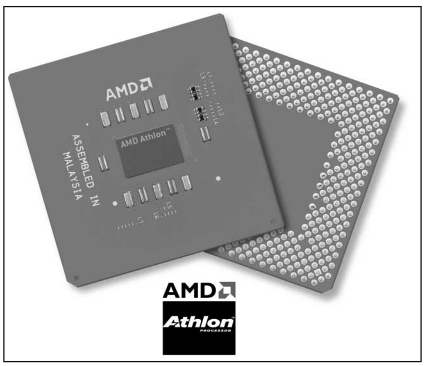

Figure 3-7 The AMD Athlon competes with the Intel Pentium III

A+CORE 4.1

Intel's Itanium: The Next Generation Processor

The next processor scheduled to be released by Intel before this book goes to print is the Itanium, Intel's first 64-bit processor for microcomputers. Recall from Chapter 1 that earlier computers always operated in real mode, which used a 16-bit data path. Later, protected mode was introduced, which uses a 32-bit data path. Almost all applications written today use 32-bit protected mode because all CPUs manufactured today for microcomputers use a 32-bit data path. The Itanium will change all that. To take full advantage of the Itanium's power, software developers must redo their applications to use 64-bit processing, and operating systems must be written to use 64-bit data transfers. Microsoft is expected to provide a 64-bit version of Windows 2000 when Itanium becomes available. Intel has promised that the Itanium will provide backward-compatibility with older 32-bit applications.

CPUs That Use RISC Technology

In addition to CPUs becoming faster and using a wider data path, another trend in chip design is the increased use of **RISC (reduced instruction set computer)** technology. RISC chips are challenging the monopoly in the chip market held by CISC (complex instruction set computer) chips. (CISC is the name given to traditional chip design.) The difference between the RISC and CISC technologies is the number of instructions (called the **instruction set**) contained directly on the CPU chip itself. With RISC technology, the CPU is limited to a very few instructions that can execute in a single clock cycle. One advantage that RISC chips have over CISC chips is that, because they have only a small number of operating instructions to perform, they can process much faster when few complex calculations are required. This feature makes RISC chips ideal for video or telecommunications applications. They are also easier and cheaper to manufacture.

Most Intel chips use the CISC technology to maintain compatibility with older systems and software, although the Pentium II uses a combination of both technologies. The K6 by AMD use the RISC technology. Cyrix, on the other hand, has chosen to stay with CISC technology, contending that it is better than RISC. Most CPU manufacturers for high-end servers have a version of a RISC chip. Sun Microsystems has the SPARC chip, Digital Equipment Corporation (DEC) has the MIPS and Alpha, and IBM Corporation has the RS 6000.

CPU Cooling Fans

Because a CPU generates so much heat, most computer systems use a cooling fan to keep the temperature below the Intel maximum allowed limit of 185° F (see Figure 3-8). Good CPU cooling fans can maintain the temperature at 90 to 110° F. Use cooling fans to prevent system errors and to prolong the life of the CPU. The ball-bearing cooling fans last longer than other kinds.

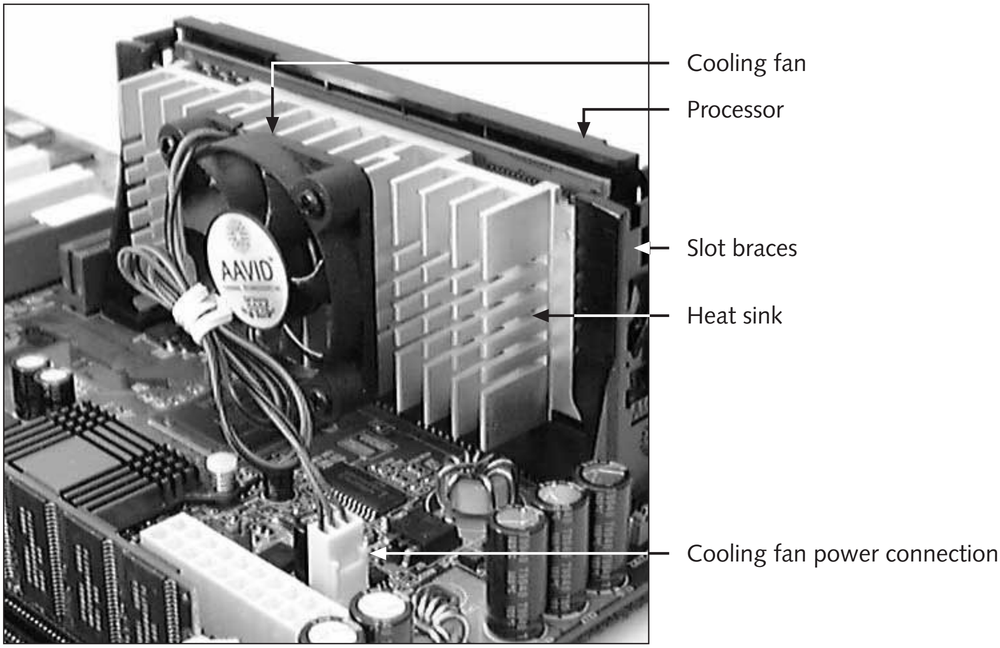

Figure 3-8 A CPU cooling fan mounts on the top or side of the CPU housing and is powered by an electrical connection to the system board

The cooling fan usually fits on top of the CPU with a wire or plastic clip. Sometimes a cream-like thermal compound is placed between the fan and the CPU. This compound draws heat from the CPU and passes it to the fan. The thermal compound transmits heat better than air and makes the connection between the fan and the CPU airtight. The fan is equipped with a power connector that connects to one of the power cables coming from the power supply.

Some newer CPUs generate so much heat that they need extra cooling. The chips might have a heat sink attached to them and a large fan attached on top of the sink or to the side of the case, blowing over the heat sink. A **heat sink** is a clip-on device that mounts on top of the CPU. Fingers or fins at the base of the heat sink pull the heat away from the CPU.

Some system boards feature a power connection for the cooling fan that sounds an alarm if the fan stops working. Because the fan is a mechanical device, it is more likely to fail than the electronic devices inside the case. To protect the expensive CPU, you can purchase a heat sensor for a few dollars. The sensor plugs into a power connection coming from the power supply and is mounted on the side of the case. It sounds an alarm when the inside of the case gets too hot.

A+ CORE 4.1

CPU Form Factors

Intel currently has five form factors used to house its processors in desktop PCs:

- **SEP (Single Edge Processor)**. The processor is not completely covered by the black plastic housing, making the circuit board visible at the bottom of the housing. The first Celeron processors used the SEP form factor in Slot 1.
- **SECC (Single Edge Contact Cartridge)**. The processor is completely covered with a black plastic housing and a heat sink and fan are attached to the housing. You can't see the circuit board or edge connector in a SECC form factor. The Pentium II and Pentium III use a SECC form factor in Slot 1. You can see the SECC in Figure 3-5.
- **SECC2 (Single Edge Contact Cartridge, version 2)**. The processor SECC2 has a heat sink and fan similarly to the SECC, but the edge connector on the processor circuit board is visible at the bottom of the housing. Pentium II and Pentium III use the SECC2 form factor.
- **PPGA (Plastic Pin Grid Array)**. The processor is housed in a square box designed to fit flat into Socket 370 (see Figure 3-9). Pins are on the underside of the flat housing, and heat sinks or fans can be attached to the top of the housing by using a thermal plate or heat spreader. Current Celeron processors use this form factor.
- **FC-PGA (Flip Chip Pin Grid Array)** This form factor looks like the PPGA form factor and uses Socket 370. Heat sinks or fans can be attached directly to the top of the package. The Pentium III uses FC-PGA as one of its two form factors.

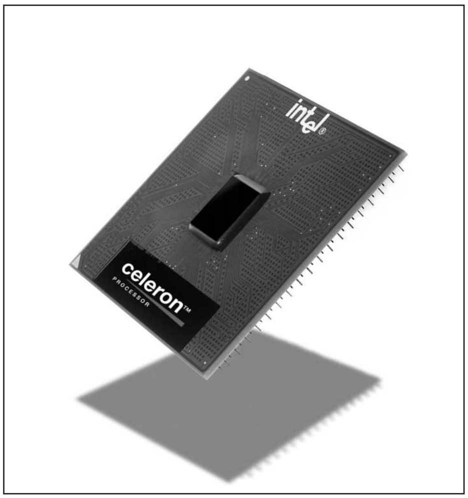

Figure 3-9 The Intel Celeron processor is housed in the PPGA form factor, which has pins on the underside that insert into Socket 370

CPU Slots and Sockets

A+CORE 1.1, 4.1, 4.3

Recall from Chapter 2 that a slot or socket is the physical connection used to connect a device (the CPU) to the system board. The type of socket or slot supplied by the system board for the processor must match that required by the processor. Table 3-7 lists several of the types of sockets and slots used by CPUs. Slots 1 and 2 are proprietary Intel slots, and Slot A is a proprietary AMD slot.

Table 3-7 CPU sockets and slots

Connector Name	Used by CPU	Number of Pins	Voltage
Socket 4	Classic Pentium 60/66	273 pins 21 × 21 PGA grid	5 V
Socket 5	Classic Pentium 75/90/100/120	320 pins 37 × 37 SPGA grid	3.3 V
Socket 6	Not used	235 pins 19 × 19 PGA grid	3.3 V
Socket 7	Pentium MMX, Fast Classic Pentium, AMD K5, AMD K6, Cyrix M	321 pins 37 × 37 SPGA grid	2.5 V to 3.3 V
Super Socket 7	AMD K6-2, AMD K6-III	321 pins 37 × 37 SPGA grid	2.5 V to 3.3 V
Socket 8	Pentium Pro	387 pins 24 × 26 SPGA grid	3.3 V
Socket 370 or PGA370 Socket	Pentium III FC-PGA, Celeron PPGA, Cyrix III	370 pins SPGA grid	1.5 V or 2 V
Slot 1 or SC242	Pentium II, Pentium III	242 pins in 2 rows Rectangular shape	2.8 V and 3.3 V
Slot A	AMD Athlon	242 pins in 2 rows Rectangular shape	1.3 V to 2.05 V
Slot 2 or SC330	Pentium II Xeon, Pentium III Xeon	330 pins in 2 rows Rectangular shape	1.5 V to 3.5 V

The 486 and earlier Pentiums used a **pin grid array (PGA)** socket where the pins were aligned in uniform rows around the socket. Later sockets use a **staggered pin grid array** (**SPGA**) where pins are staggered over the socket to squeeze more pins into a small space. Socket 7 is used on system boards that run at 66 MHz, and Super Socket 7 is used on the newer 100-MHz system boards. Super Socket 7 runs on 100-MHz system boards and supports AGP. It was designed to be used with AMD CPUs competing with the Intel Pentium II.

Socket 370 is used by two types of processors with two types of package form factors: PPGA (Plastic Pin Grid Array) and FC-PGA (Flip Chip Pin Grid Array). Both form

A+CORE 1.1, 4.1, 4.3

factors have pins on the underside of the processor that insert into the pin holes on the Socket 370 (see Figure 3-9). Socket 370 is also used by the Cyrix III and a version of the Pentium III designed for smaller computer cases called the Pentium III FC-PGA. For socket comparisons, see Figure 3-10.

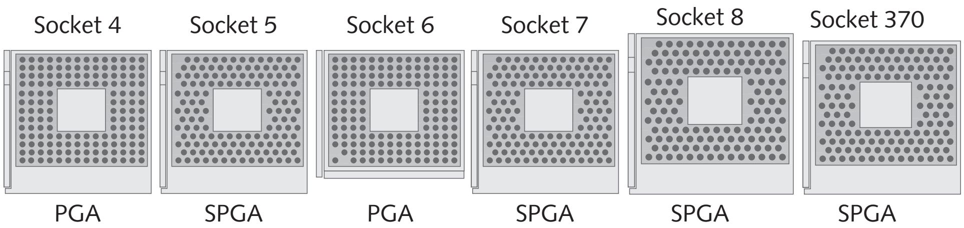

Figure 3-10 CPU sockets use either a PGA or SPGA design; rows of pins are arranged on the socket either in even rows (PGA) or staggered (SPGA)

Earlier CPU sockets, called dual inline pin package (DIPP) sockets, were rectangular with two rows of pins down each side. PGA and SPGA sockets are all square or close to it. DIPP and some PGA sockets, called **low insertion force** (**LIF**) sockets, were somewhat troublesome to install because it was difficult to apply even force when inserting them. Current CPU sockets are called **zero insertion force** (**ZIF**) sockets and have a small lever on the side of the socket that lifts the CPU up and out of the socket. Push the lever down and the CPU moves into its pin connectors with equal force over the entire housing. The heat sink or fan clips to the top of the CPU. With this method, you can more easily remove the CPU and replace it with another if necessary.

Slot 1, Slot A, and Slot 2 are all designed to accommodate processors using the SEP or SECC housings that stand on their end much like an expansion card. The CPU is secured in the slot with clips on each side of the slot. You can attach a heat sink or cooling fan to the side of the CPU case. The Pentium II and Pentium III use Slot 1, and the Xeon versions of these processors use the longer Slot 2. AMD processors use Slot A.

The Celeron processor uses the PPGA form factor and Socket 370. Some system boards that have a Slot 1 can accommodate the Celeron processor by using a riser CPU card (see Figure 3-11). The riser card inserts into Slot 1, and the Celeron processor inserts into Socket 370 on the riser card. This feature allows you to upgrade an older Pentium II system to the faster Celeron.

The Pentium III processor uses two types of form factors: the SECC 2 and the FC-PGA. The SECC 2 is inserted into Slot 1, and the FC-PGA uses Socket 370. However, some older system boards that have a Socket 370 are not designed to support the Pentium III FC-PGA.

A+ CORE 1.1, 4.1, 4.3

The Celeron processor uses 2.00 volts, and the Pentium III FC-PGA uses either 1.60 or 1.65 volts. A system board built only to support the Celeron may not recognize the need to step down the voltage for the Pentium III FC-PGA processor. This overvoltage can damage the Pentium III. Always consult the system board documentation to know what processors the board can support.

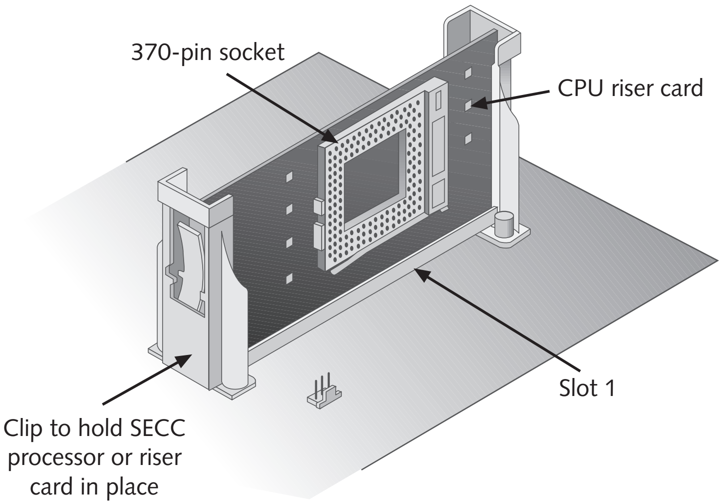

Figure 3-11 A riser card can be used to install a Celeron processor into a system board with slot 1

CPU Voltage Regulator

As you can see from Table 3-7, different CPUs require different amounts of voltage on the system board. Some CPUs require one voltage amount for external operations and another amount for internal operations. Those that require two different voltages are called **dual-voltage CPUs**. The others are called **single-voltage CPUs**. A CPU voltage regulator controls the amount of voltage on the system board. Some CPUs require that you set the jumpers on the system board to control the voltage, and other CPUs automatically control the voltage without your involvement.

Figure 3-12 shows sample documentation of what jumper settings to use for various CPU voltage selections on a particular system board that require jumpers to be set. Notice that two jumpers called JP16 located near Socket 7 on the board accomplish the voltage selection. For single voltage, for the Pentium, Cyrix 6x86, or AMD K5, both jumpers are open. Dual voltage used by the Pentium MMX, Cyrix M2, and AMD K6 is selected by opening or closing the two jumpers according to the diagram. Follow the recommendations for your CPU when selecting the voltages from the chart.

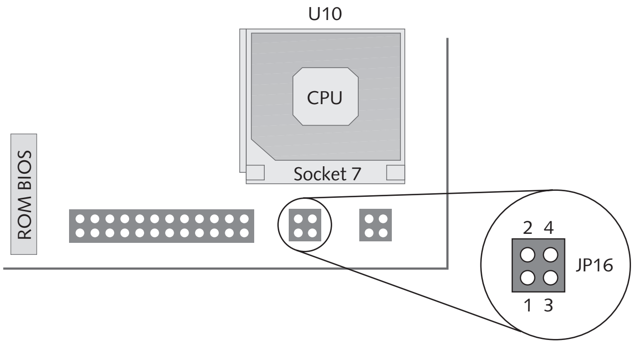

CPU TYPE	JP16	CPU Voltage CORE	CPU Voltage I/O
Single Voltage INTEL P54C/CQS/CT Cyrix 6x86 AMD K5	2 4 1 3 open	3.5V	3.5V
Dual Voltage INTEL P55C/MMX Cyrix 6x86L/M2 AMD K6	2 4 1 3 open	2.8V	3.4V
	2 4 1 3 1-2 closed, 3-4 open	2.9V	3.4V
	2 4 1 3 1-2 open, 3-4 closed	3.2V	3.4V

Figure 3-12 CPU voltage regulator can be configured using jumpers on the system board to apply the correct voltage to the CPU

THE CHIP SET

A chip set is a set of chips on the system board that collectively controls the memory cache, external buses, and some peripherals. Intel makes the most popular chip sets, which are listed in Table 3-8.

The Intel 440BX chip set is the first PC chip set to offer a memory bus that runs at 100 MHz, allowing a Pentium II running at 350 MHz or 400 MHz to reach its full potential for performance in desktop PCs. Before this, the memory bus slowed the CPU speed. The 440BX chip set is also the first chip set to use the mobile version of the Pentium II processor for notebooks. Often you see this chip set advertised with "AGP" in the name, as in the Intel 440BX AGP chip set. The 440GX chip set is an evolution of the 440BX.

Table 3-8 The Intel chip set family

Common Name	Model Number	Comments
Intel i800 Series	840	Designed for multiprocessor systems using Pentium II Xeon or Pentium III Xeon processors
	820	Designed for Pentium II and Pentium III systems
	810	First Intel chip set to eliminate the PCI bus as the main device interconnection
Orion	450GX, KX	Supports Pentium Pro (includes support for multiprocessors)
	450NX	Designed for servers with multiple Pentium II or Pentium II Xeon processors
Natoma	440FX	Supports Pentium Pro and Pentium II (Discontinued in January, 1999)
	440BX	Designed for servers and workstations (Pentium II and III)
	440GX	Designed for servers and workstations using the Pentium II Xeon and Pentium III Xeon
	440ZX	Designed for entry-level PCs using Pentium II
	440LX	Designed for Celeron processors
	440MX	Designed for notebook computers (M = mobile)
	440EX	Designed for smaller system boards such as the mini-ATX
Triton III	430VX	Value chip set, supports SDRAM
	430MX	Used for notebooks (M = mobile)
	430TX	Supports SDRAM, ultra DMA; replaced the VX and MX
Triton II	430HX	High performance, supports dual processors
Triton I	430FX	The oldest chip set, no longer produced

A+ CORE 4.3

The 400 series of Intel chip sets uses the PCI bus as the interconnection between slower buses and the system bus. How the PCI bus does this is covered later in the chapter. The Intel i800 series of chip sets introduced a new way for I/O buses to relate to the faster system bus and ultimately to the CPU. With the i800 series, the interconnection between buses is done using a hub interface architecture, whereby all I/O buses connect to a hub, which connects to the system bus. This hub is called Hub Interface, and the architecture is called the Accelerated Hub Architecture (see Figure 3-13). The fast end of the hub, which contains the graphics and memory controller (GMCH), connects to the system bus and is called the hub's North Bridge. The slower end of the hub, called the South Bridge, contains the I/O Controller Hub (ICH). All I/O devices except display and memory connect to the hub by using the slower South Bridge.

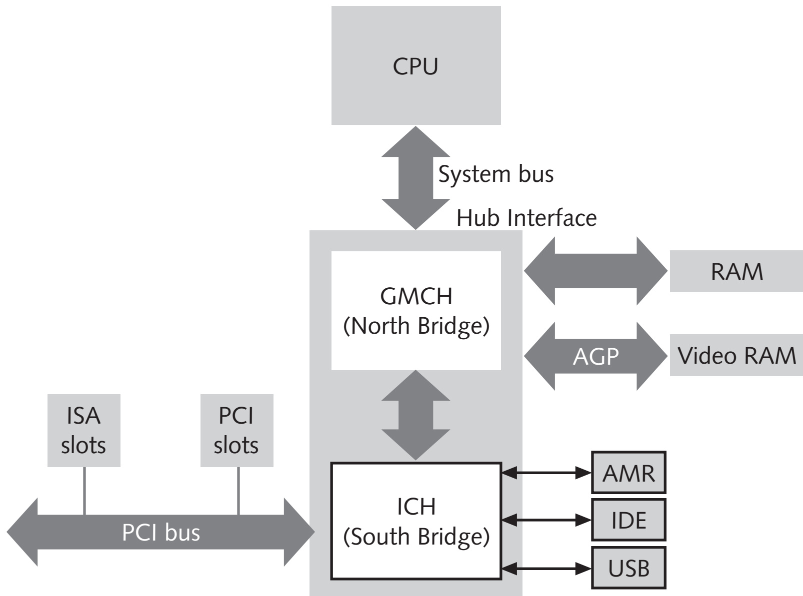

Figure 3-13 Using Intel 800 series Accelerated Hub Architecture, a Hub Interface is used to connect slower I/O buses to the system bus

A+ CORE 4.3

Chip sets are manufactured by the following companies:

- Intel Corporation
- Cyrix Corporation
- Silicon Integrated Systems Corp. (known as SiS)
- ALi, Inc.
- Standard Microsystems Corp.
- United Microelectronics Corp.
- VIA Technology, Inc. combined with AMD, Inc.
- VLSI Technology, Inc.

Chip Sets that Compete with Intel

SiS uses a similar faster North Bridge and slower South Bridge approach to managing slower I/O buses interconnecting with the faster system bus. The SiS620 chip set includes a digital video interface for digital flat panel display screens, and supports a 100 MHz system bus and an advanced hard drive interface called Ultra DMA (discussed in later chapters).

The Aladdin V chip set from ALi supports Socket 7 processors and the 100-MHz system bus speed. VIA in combination with AMD has the Apollo MVP3 chip set, which supports AGP, a 100-MHz bus speed, and Socket 7.

A+CORE 4.3

Currently, Intel dominates the chip set market for several reasons. The major advantage that Intel has over other chip set manufacturers is that they know more about their Intel CPUs than anyone else, and the chip sets are therefore more compatible with the Pentium family of CPUs. Intel's investment in research and development has let their engineers invent the PCI bus, the universal serial bus, the accelerated graphics port (AGP), and more recently the Accelerated Hub Architecture.

ROM BIOS

A+CORE 1.1, 1.8

Recall that there is one ROM chip on the system board that contains BIOS, which manages the startup process (startup BIOS) and many basic I/O functions of the system (system BIOS). Phoenix Software, Award Software, and American Megatrends, Inc. (AMI) write the most well-known and dependable ROM code for PCs. When selecting a PC clone, make sure to check who wrote the ROM BIOS code. If you select code written by one of these companies, your ROM BIOS will be compatible with most software.

An easy way to identify the name of the BIOS manufacturer without having to remove the case cover is to watch the boot process. The name of the BIOS manufacturer appears at the beginning of the boot process. You can also look for identifying information written on top of the chip. This ROM BIOS chip is easy to spot because it is larger than most chips and often has a shiny plastic label on it. On the label is the manufacturer's name, the date of manufacture, and the serial number of the chip. This information is important when you are trying to precisely identify the chip, such as when you're selecting the correct upgrade for the chip.

In the past, if the ROM BIOS needed upgrading—either because of new hardware or software added to the system or because the BIOS was causing errors, this meant exchanging the chip. The chip is usually socketed in, not soldered, for easy exchange. Recall from Chapter 1 that a newer kind of ROM, called Flash ROM, is now available that allows upgraded versions of the BIOS to be written to it without having to physically replace the chip.

You need to know the following about your BIOS:

- Does the BIOS support Plug and Play?
- Does the BIOS support large hard drives?
- Is the BIOS chip a Flash ROM chip?

This one ROM BIOS chip on the system board contains only a portion of the total BIOS code needed to interface with all the hardware components in the system. Understanding that BIOS programs can come from several sources helps in solving memory problems and other problems that arise from resource conflicts that will be considered in future chapters.

The Total BIOS in Your System

A+CORE 1.1, 1.8

Some expansion cards, such as a network interface card (NIC) or a video/graphics card, also have ROM chips on them containing BIOS code. The operating system uses the programs stored on these ROM chips to communicate with the peripheral devices. During the boot

A+ CORE 1.1, 1.8

process, the expansion card tells the startup program how many memory addresses that it requires to access its ROM code. For protected mode firmware, any memory addresses will do, but, for older legacy cards using real mode, the BIOS must be assigned addresses in base memory or the upper memory area between 640K and 1024K. The ROM code from these boards becomes part of the total BIOS that the OS uses to communicate with peripherals. Problems referred to as hardware configuration conflicts can occur if two legacy boards request the same addresses in upper memory. We will address these hardware conflicts and offer possible solutions when we study managing memory in Chapter 4.

Figure 3-14 shows how the programming code from various ROM BIOS chips can be mapped onto the memory addresses managed by the CPU. The areas of upper memory are labeled the F range and the C range. In hex notation, upper memory addresses are numbered A0000 to FFFFF. Because of these hex numbers, the divisions of upper memory are often referred to as the A range, B range, C range, and so on, up to F range.

Recall from Chapter 2 that memory is viewed logically as a series of memory addresses that can be assigned to physical memory devices, such as a SIMM on the system board, a ROM BIOS chip on the system board, or a ROM chip on a network card or modem card. After booting is complete, most if not all of the BIOS on the system has declared that it exists and has requested memory addresses. In Figure 3-14, each memory device has been assigned a different address in base, upper, or extended memory.

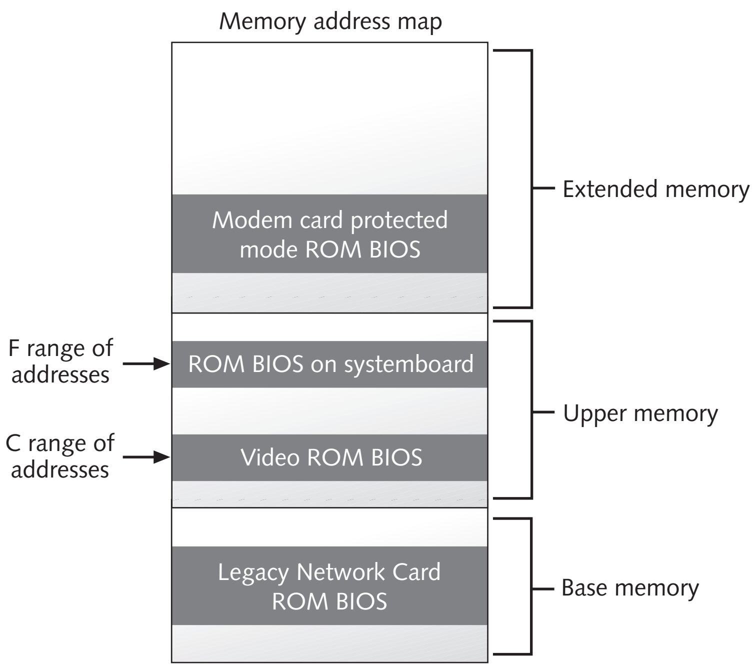

Figure 3-14 The total ROM BIOS programs in a system can be assigned memory addresses in base, upper, and extended memory

A+ CORE 1.1, 1.8

Remember from Chapter 2 that if the programming code from the ROM BIOS chips is also copied into RAM, this is called shadowing ROM or sometimes Shadow RAM. These terms indicate that RAM is shadowing ROM code. In the setup of your computer, you usually have the choice of whether to shadow System BIOS. For DOS and Windows 9x, accept the default setting for this option.

Plug and Play BIOS

Recall from Chapter 1 that **Plug and Play** (**PnP**) is a term that applies to both the Windows 9x OS and to some ROM BIOS. It means that rather than having you reset DIP switches and jumpers, the OS and/or the BIOS automatically configures hardware devices to reduce or eliminate conflicting requests for such system resources as I/O addresses, IRQs, DMA channels, or upper memory addresses. Windows 9x Plug and Play assigns these resources to a device only if the device allows it. For example, if a legacy sound card requires a certain group of upper memory addresses that are hard coded into its on-board BIOS, there's nothing that Windows 9x Plug and Play can do about that. (Hard coded is computer jargon for something being coded so that it cannot be changed.) Plug and Play simply tries to work around the problem as best it can. If two non-Plug and Play hardware devices require the same resource and their BIOS does not provide for accepting a substitute, these two devices cannot coexist on the same PC.

Newer devices that are Plug and Play compliant are more cooperative. At startup, they simply request to work and then wait for the OS to assign the resources they need. Windows 9x and Windows 2000 try to do that whether or not the system BIOS is Plug and Play BIOS. Plug and Play BIOS does some of the up-front work for Windows, the way that an efficient secretary organizes a boss's work for the day. At startup, it's the Startup BIOS that examines the hardware devices present, takes inventory, and then loads the OS. Part of the job of Plug and Play BIOS is to collect information about the devices and the resources they require and later to work with Windows 9x or Windows 2000 to assign the resources.

ESCD (extended system configuration data) Plug and Play BIOS goes even further, creating a list of all the things you have done manually to the configuration that Plug and Play does not do on its own. This ESCD list is written to the BIOS chip so that the next time you boot, the Startup BIOS can faithfully relate that information to Windows. The BIOS chip for ESCD BIOS is a special RAM chip called Permanent RAM, or PRAM, that can hold data written to it without the benefit of a battery, which the CMOS setup chip requires.

Most ROM BIOS chips made after the end of 1994 are Plug and Play. Windows 9x and Windows 2000 can use most—but not all—of their Plug and Play abilities without Plug and Play BIOS. If you are buying a new PC, accept nothing less than Plug and Play BIOS. As more and more devices become Plug and Play compliant, the time will come when installing a new device on a PC will be just as error-free and easy to do as it is on a Mac; Apple for years has known about and used the same concepts as are used in Plug and Play.

When BIOS is Incompatible with Hardware or Software

A+ CORE 2.1, 4.4, 1.8

BIOS is a hybrid of two worlds. It's technically both hardware and software—it's really the intersection point of the two—and must communicate with both well, as shown in Figure 3-15. When hardware and software change, BIOS might need to change too. In the past, most users upgraded BIOS because new hardware was incompatible with it. Sometimes, however, you need to upgrade BIOS to accommodate new software, such as Plug and Play.

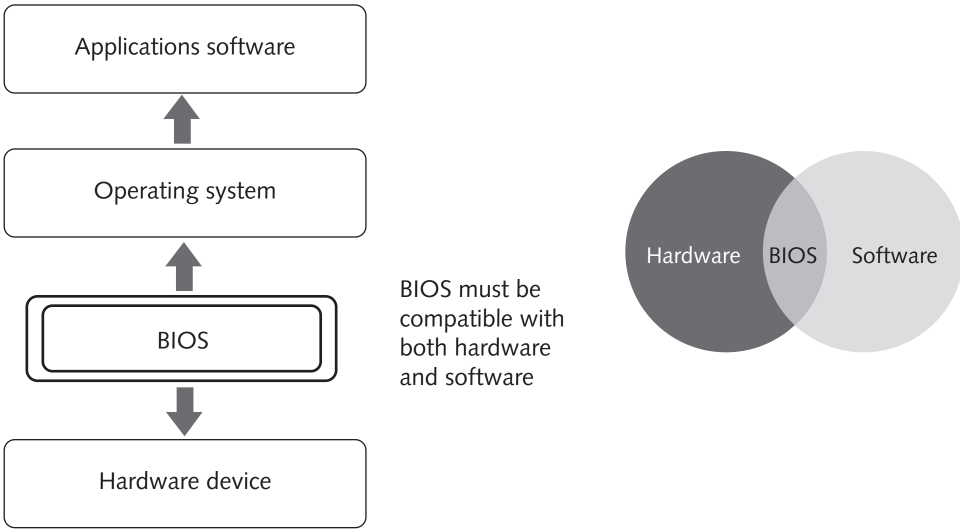

Figure 3-15 BIOS can serve as the hardware/software interface

Many years ago, when a new device became available, such as the 3½ inch floppy disk drive, your PC sometimes could not use the new device until you upgraded the BIOS. You did that by replacing the old BIOS chip with a new chip that supported the new device. Now, however, it's much easier. First, remember that most of today's new devices are not supported by the System BIOS at all, but by device drivers, which are software programs installed on the hard drive as an add-on part of the OS. But, if some new feature does require an upgrade to BIOS, you can do that with Flash ROM. Installing a larger hard drive is an example of a hardware upgrade that might require a BIOS upgrade because it is incompatible with the existing BIOS. Older BIOS supports only those hard drives with a 504-MB capacity. If you have this problem—large drive, old BIOS—you can solve it in one of two ways. Either upgrade BIOS or use special software designed to get around the problem. Often the device manufacturer supplies the software.

Flash ROM

Technically speaking, Flash ROM is called EEPROM (electronically erasable programmable read-only memory), which means you can change the programming on the chip through software on your PC. The updated programming will be retained—even when you turn off

A+CORE 2.1, 4.4, 1.8

your PC for long periods of time—until you change it again. Flash ROM allows you to upgrade system BIOS without having to replace the ROM chip.

As more devices become Plug and Play compliant, Plug and Play BIOS will become more sophisticated. Additionally, makers of BIOS code are likely to change BIOS frequently because it is so easy for them to provide the upgrade on the Internet. You can get upgraded BIOS code from manufacturers' web sites or disks or from third-party BIOS resellers' Web sites or disks.

Figure 3-16 shows a sample Web site for Flash ROM BIOS upgrades. See the web site of the manufacturer of your BIOS or of the system board manufacturer for more information.

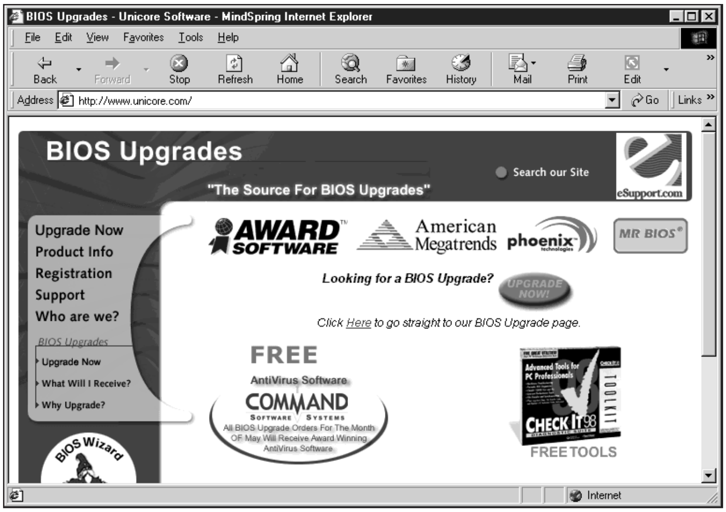

Figure 3-16 Flash ROM BIOS upgrades for most BIOS manufacturers can be downloaded from *www.unicore.com*

To upgrade Flash ROM, follow the directions that came with your system board and the upgrade software itself. Generally, you perform these tasks:

- Set a jumper on the system board telling the BIOS to expect an upgrade
- Copy the upgrade BIOS software to a bootable disk
- Boot from the disk and follow the menu options to upgrade the BIOS
- Set the jumper back to its original setting, reboot the system, and verify that all is working

Be *very careful* that you upgrade the BIOS with the correct upgrade and that you follow manufacturer instructions correctly. Upgrading with the wrong file could make your system BIOS totally useless. If you're not sure that you're using the correct upgrade, *don't guess*. Check with the technical support for your BIOS before moving forward. Before you call technical support, have the information that is written on the BIOS chip label available.

3

RAM (Random Access Memory)

A+ CORE 4.2

Chapter 4 discusses how to manage RAM, but for now we present the essentials of where and what RAM is and how it is used. In older machines, RAM existed as individual chips socketed to the system board in banks or rows of nine chips each. Each bank held one byte by storing one bit in each chip, with the ninth chip holding a **parity** bit (see Figure 3-17). On older PCs the parity chip was separated a little from the other eight chips. (Parity, discussed below in more detail, refers to an error-checking procedure whereby either every byte has an even number of ones or every byte has an odd number of ones. The use of a parity bit means that every byte occupies nine rather than eight bits.)

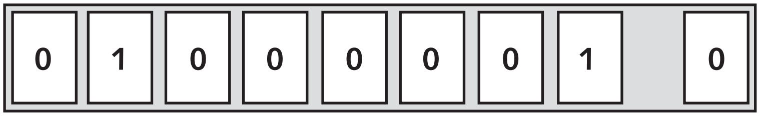

Figure 3-17 Eight chips and a parity chip represent the letter A in ASCII with even parity

Dynamic Memory

Recall that there are two types of RAM: **dynamic RAM** (**DRAM**) and **static RAM** (**SRAM**). Dynamic RAM chips hold data for a very short time; static RAM chips hold data until the power is turned off. Because DRAM is much less expensive than SRAM, most of the RAM on the system board is DRAM. DRAM comes in three types: parity, nonparity, or an altogether new method of error checking called ECC (error correcting code) that cannot only detect an error but also correct it. More about ECC in Chapter 4.

Parity is a method of testing the integrity of the bits stored in RAM or some secondary medium, or testing the integrity of bits sent over a communications device. When data is written to RAM, the computer calculates how many ON bits (binary 1) there are in the 8 bits of a byte. If the computer uses odd parity, it makes the ninth or parity bit either a 1 or a 0 to make the number of 1s in the 9 bits odd. Using even parity, the computer makes the parity bit a 1 or 0 to make the number of 1s in the 9 bits even.

Later, when the byte is read back, the computer checks the odd or even state. If the number of bits is not an odd number for odd parity or an even number for even parity, a **parity error** occurs. A parity error always causes the system to halt. On the screen you see the error message “Parity Error 1” or “Parity Error 2” or a similar error message about parity. Parity Error 1 is a parity error on the system board; Parity Error 2 is a parity error on a memory expansion board. Parity errors can be caused by RAM chips that have become

A+ CORE 4.2

undependable and that are unable to hold data reliably. Sometimes this happens when the chips overheat or power falters.

Recall that later, computers were made to hold RAM on a group of chips stored in a single physical unit called a **SIMM** (single inline memory module). A SIMM is a miniboard that stores an entire bank or banks of RAM. A SIMM can have several chips with 30 or 72 pins on the edge connector of the tiny board. RAM is then upgraded or changed by unplugging and plugging in SIMMs, which are much easier to work with than single chips. You will learn to upgrade memory using SIMMs in Chapter 4. RAM chips or SIMMs are located either on the system board or on memory expansion cards. SIMMs hold from 8 MB to 64 MB of RAM on one board.

All new system boards today use **DIMMs** (dual inline memory module), which have 168 pins on the edge connector of the board. A DIMM can hold from 8 MB to 256 MB of RAM on a single board. Figure 3-18 shows the two kinds of SIMMs and a DIMM module.

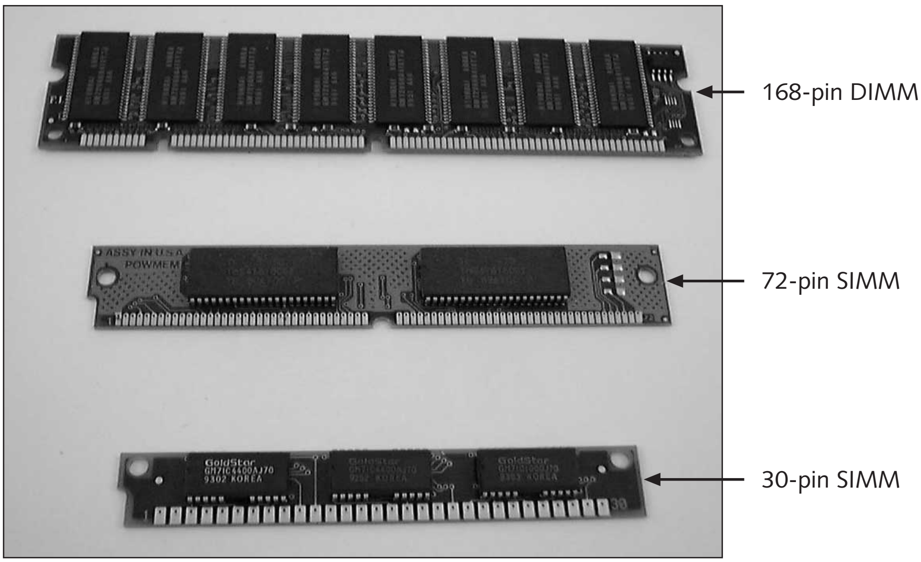

Figure 3-18 Types of RAM modules

Most system boards today use 168-pin DIMMs. However, memory can be managed using several technologies that involve how memory is accessed, how timing the access is managed, and how the system board and the CPU relate to the memory modules. The more prevalent memory technologies (and some variations of each) used by the industry are listed in Table 3-9. Each of these technologies will be discussed in detail in Chapter 4. For now, know that the technology used by the memory modules must match the technology supported by the system board.

A+CORE 4.2

Table 3-9 DRAM memory technologies

Technology	Description
Conventional	Used with earlier PCs but currently not available
Fast Page Memory (FPM)	Improved access time over conventional memory. FPM may still be seen today.
Extended Data Out (EDO)	Refined version of FPM that speeds up access time. Still seen on older system boards.
Burst EDO (BEDO)	Refined version of EDO that significantly improved access time over EDO. BEDO is seldom used today because Intel chose not to support it.
Synchronous DRAM (SDRAM)	SDRAM runs in sync with the system clock and is rated by clock speed, whereas other types of memory run independently of (and slower than) the system clock.
Rambus DRAM (RDRAM)	RDRAM uses a faster memory bus (up to 800 MHz), but only a 16-bit data path. The Intel Itanium processor will use RDRAM.
Double Data Rate (DDR) SDRAM	A faster version of SDRAM that can run at 200 MHz

Regardless of the type, dynamic RAM chips do not hold their data very long and must be refreshed about every 3.86 milliseconds. To **refresh** RAM means that the computer must rewrite the data to the chip. Refreshing RAM is done by the DMA (dynamic memory access) chip (discussed later in this chapter) or sometimes by circuitry on the system board other than the DMA chip.

Static Cache Memory

Recall that there are two kinds of static memory. L1 is contained on the CPU microchip, and L2 cache is external to the chip. L2 cache is housed either on the system board or, for newer CPUs, inside the CPU case. Look back at Figure 3-2 to see an example of an older CPU (the Pentium MMX) system board that has a memory cache slot. In Chapter 4 you will learn how to install extra L2 memory cache on one of these older boards.

Buses and Expansion Slots

As cities grow, so do their transportation systems. Small villages have only simple, two-lane roads, but large cities have two-lane and four-lane roads, and major freeways, each with its own set of traffic laws, including minimum and maximum speeds, access methods, and protocols. As microcomputer systems have evolved, so too have their "transportation" systems. The earliest PC had only a single and simple bus. Today's PCs have four or five buses, each with different speeds, access methods, and protocols. As you have seen, backward-compatibility dictates that older buses still be supported on a

system board, even when faster, better buses exist. All this makes for a maze of many buses on a system board.

Bus Evolution

Just as a city's road system improves to increase the speed and number of lanes of traffic, buses have evolved around these similar issues, data path and speed. Cars on a freeway travel at a continuous or constant speed, but traffic on a computer's CPU or bus travels in a digital (on and off) manner rather than in an analog (continuous) manner. The system clock, run by a crystal on the system board, occupies one line of a bus and keeps the beat for components. Do something. Stop. Do something. Stop. Do something. Stop. With each beat, called a clock cycle, something can happen. Everything stops between beats, waiting for the next beat. The CPU is listening to this beat and working on these clock cycles. If another component on the system board also works by the beat or clock cycle, then it is said to be synchronized with the CPU. For example, earlier in the chapter, it was said that the backside bus of the Pentium II worked at half the speed of the CPU. This means that the CPU is doing something on each clock cycle, but the backside bus is doing something on every other clock cycle.

Some components don't attempt to keep in sync with the CPU, even to work at a half or a third of clock cycles. These components are said to be working asynchronously with the CPU. They might be working at a rate determined by the system clock or by another crystal on or off the system board. Either way, the frequency will be much slower than the CPU and not in sync with it. If the CPU requests something from one of these devices, and the device is not ready, it will issue wait states to the CPU until it can catch up.

A+ CORE 4.3

Devices attached to the 8-bit or 16-bit ISA bus are an example of these slower devices. The 16-bit ISA bus works at a rate of 8.33 MHz, compared to memory bus speeds of 66 MHz to 200 MHz. Buses that work in sync with the CPU and the system clock are called **local buses** (sometimes called **system buses**). Buses that work asynchronously with the CPU at a much slower rate are called **expansion buses**. The memory bus is a local bus, and the ISA bus is an expansion bus.

Table 3-10 Buses listed by throughput in MB/sec (megabytes per second) or Mbps (megabits per second)

Bus	Bus Type	Data Path in Bits	Address Lines	Bus Speed in MHz	Throughput
FireWire	Local I/O or expansion	1	Addresses are sent serially	NA	Up to 1.2 GB/sec
Memory bus	Local	64	32	66, 75, 100 …	Up to 1 GB/sec
AGP	Local video	32	NA	66, 75, 100 …	Up to 528 MB/sec
PCI	Local I/O	32	32	33, 66	Up to 264 MB/sec

A+CORE 4.3

Table 3-10 Buses listed by throughput in MB/sec (megabytes per second) or Mbps (megabits per second) (continued)

Bus	Bus Type	Data Path in Bits	Address Lines	Bus Speed in MHz	Throughput
VESA or VL Bus	Local video or expansion	32	32	Up to 33	Up to 250 MB/sec
MCA	Expansion	32	32	12	Up to 40 MB/sec
EISA	Expansion	32	32	12	Up to 32 MB/sec
16-bit ISA	Expansion	16	24	8.33	8 MB/sec
8-bit ISA	Expansion	8	20	4.77	1MB/sec
USB	Expansion	1	Addresses are sent serially	3	1.5 or 12 Mbps

Why So Many Buses?

A+CORE 4.3

When the first PCs were introduced in the early 1980s, there was only one bus on the system board, called the system bus, which ran at the same speed as the CPU (4.77 MHz). Everything on the system board working with the CPU or the bus would simply keep the same beat, following the pulses of the one system clock. (This first bus is now called the 8-bit ISA bus.) Things today are not so simple. With the speeds of different hardware components evolving at different rates, a single speed for all components is no longer practical. The CPU works at one speed, the bus connecting the CPU to memory is working at a slower speed, and the bus communicating with I/O devices must work at an even slower speed. As manufacturers attempt to improve performance, new buses are invented to accommodate the characteristics of particular devices. In fact, there might be as many as five or six different buses working at different speeds on the same system board. Each bus will have a set speed that all components connected to it will work at. There are components that convert data moving from bus to bus to the speed of the new bus.

Table 3-10 lists the system-board buses, some outdated and some in use today, ordered from fastest to slowest. Historically, the 8-bit ISA (Industry Standard Architecture) bus came first, and was later revised to the 16-bit ISA bus to keep up with the demand for wider data path sizes. Then, in 1987, IBM introduced the first 32-bit bus, the MCA (Microchannel Architecture) bus, and competitors followed with the 32-bit EISA (Extended Industry Standard Architecture) bus. Because these buses are not synchronized with the CPU, they are all expansion buses. Of these buses, the only one still in use today is the 16-bit ISA. A relatively new expansion bus is the universal serial bus (USB), which targets slow I/O devices such as the mouse, digital camera, and scanner. Its advantage is that USB devices are easily installed and configured.

A+ CORE 4.3

A local bus is synchronized with the CPU. In the sense that a local bus is a bus that is close to or "local to" the CPU, there is only one "true" local bus, the memory bus or system bus, which connects directly to the CPU. All other buses must connect to it to get to the CPU. A local I/O bus is a bus designed to support fast I/O devices such as video and hard drives; it runs synchronized with the system clock, which means that it is also synchronized with the CPU. Local I/O buses did not always exist on a PC but were created as the need arose for a bus that was synchronized with the system clock, and was not as fast as the memory bus, but was faster than an expansion bus. The evolution of local I/O buses includes earlier proprietary designs, the VESA bus, the PCI bus, and the newer AGP bus. Of these, only the PCI and AGP bus are still sold. The FireWire bus is the latest local I/O bus and is not readily available as yet. It can work either synchronously or asynchronously and so is classified as either a local or expansion bus. The VESA bus could also be set to work either way.

What a Bus Does

Look on the bottom of the system board and you will see a maze of circuits that make up a bus. These embedded wires are carrying four kinds of cargo:

- **Electrical power**. Chips on the system board require power to function. These chips tap into a bus's power lines and draw what they need.
- **Control signals**. Some of the wires on a bus carry control signals that coordinate all the activity.
- **Memory addresses**. Memory addresses are passed from one component to another as these components tell each other where to access data or instructions. The number of wires that make up the memory address lines of the bus determines how many bits can be used for a memory address. The number of wires thus limits the amount of memory the bus can address.
- **Data**. Data is passed over a bus in a group of wires, just as the memory addresses are. The number of lines in the bus used to pass data determines how much data can be passed in parallel at one time. The number of lines depends on the type of CPU and determines the number of bits in the data path. (Remember that a data path is that part of the bus on which the data travels and can be 8, 16, 32, 64 or more bits wide.)

Most often when comparing buses users focus on the width of the data path and the overall bus speed. But you also should consider the type of expansion slot the bus allows. The number of fingers on the edge connector of the expansion card and the length of the edge connector are determined by the bus that controls that expansion slot. Various bus connections are shown in Figure 3-19.

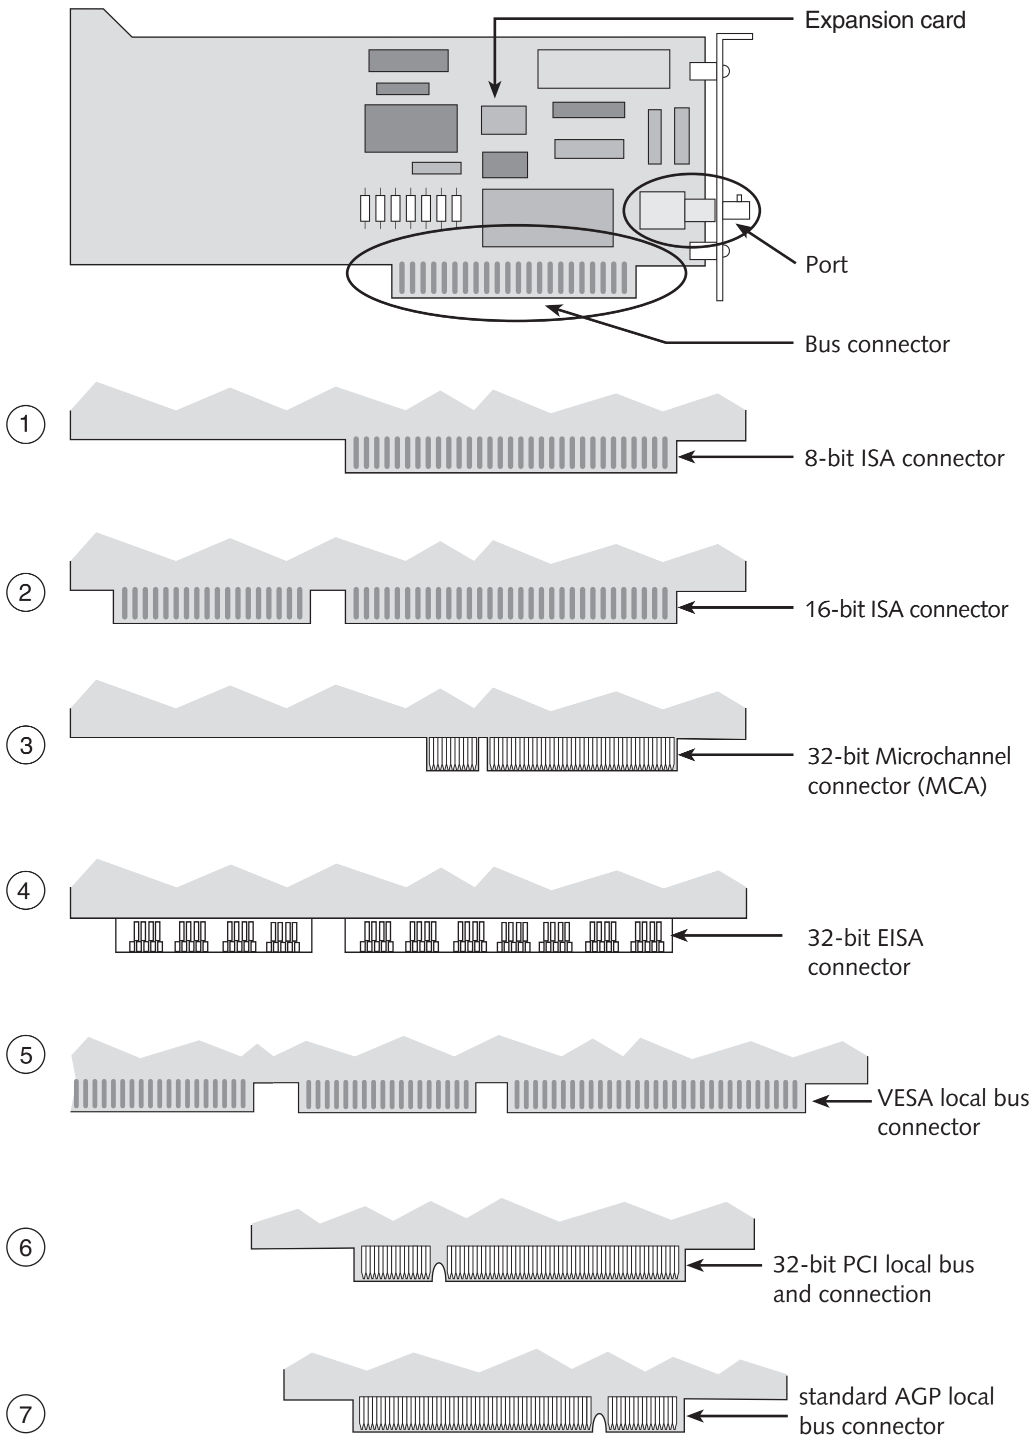

Figure 3-19 Seven bus connections on expansion cards

The ISA Bus

A+CORE 4.3

Used on the first IBM 8088 PCs in the early 1980s, the **ISA bus** had an 8-bit data path. Later, IBM revised the ISA bus to have a 16-bit path. The IBM AT personal computer used this bus and the 80286 chip, which is why the 16-bit bus is sometimes called the AT bus. IBM wanted this bus to be backward-compatible with the older 8-bit ISA bus so that the older 8-bit circuit boards would fit into the newer AT computers. To maintain compatibility, IBM kept the old 62-line slot connector and added another slot connector beside it to provide the extra 8 bits. Slots with both connectors are called 16-bit slots. A new system board today usually has at least one 16-bit slot, which can be used by either an 8-bit or 16-bit ISA card.

Microchannel Architecture (MCA) Bus

With the introduction of the line of PS/2 computers in 1987, IBM introduced the first 32-bit bus for personal computers, the **Microchannel Architecture (MCA) bus**. IBM did not intend the MCA bus to be compatible with ISA buses. Circuit boards used in older IBM computers could not be used in the PS/2 line. (The PS/2 Models 25 and 30 still included the older ISA bus in order to support legacy cards.)

IBM chose to patent the bus so that other companies could not economically manufacture and market it. IBM intended to control a subset of the bus market with MCA. In response, Compaq and eight other companies (called the "Gang of Nine") joined to design and build a competing 32-bit bus, the EISA bus.

The EISA Bus

Designed to compete with the MCA bus, the **EISA (Extended ISA) bus** (pronounced "ease-sa") has a 32-bit data path. The bus is compatible with older ISA buses so that expansion boards having 8-bit or 16-bit data paths work on the EISA bus. The speed of the EISA bus is about 20 MHz. To accommodate a 16-bit or 8-bit ISA circuit board, the 32-bit EISA has two slots that have the same width as 16-bit ISA slots. However, the EISA bus slots are deeper than 16-bit slots. All 32-bit circuit boards have longer fingers on the edge connectors that go deep into the EISA slot connecting to the 32-bit pins. A 16-bit circuit board reaches only partway down the slot connecting at a shallower level only to the 16-bit pins.

Universal Serial Bus

A+CORE 4.3, 1.3, 1.4

A relatively new I/O bus is the **universal serial bus** or **USB**, originally created by a seven-member consortium including Compaq, Digital Equipment, IBM, Intel, Microsoft, NEC, and Northern Telecom. It is designed to make the installation of slow peripheral devices as effortless as possible. USB is much faster than regular serial ports and much easier to manage, eliminating the need to manually resolve resource conflicts, since the host controller only uses one set of resources for all devices. It is expected that USB will ultimately replace both serial and parallel ports as the technology matures and more devices are built to use USB.

One or two USB ports are found on most new system boards today (see Figure 3-20), and older system boards that don't have USB ports can be upgraded by adding a PCI-to-USB controller card in a PCI slot to provide a USB port. USB allows for two speeds, 1.5 Mb per second and 12 Mb per second, and works well for slow I/O devices.

A+ CORE 4.3, 1.3, 1.4

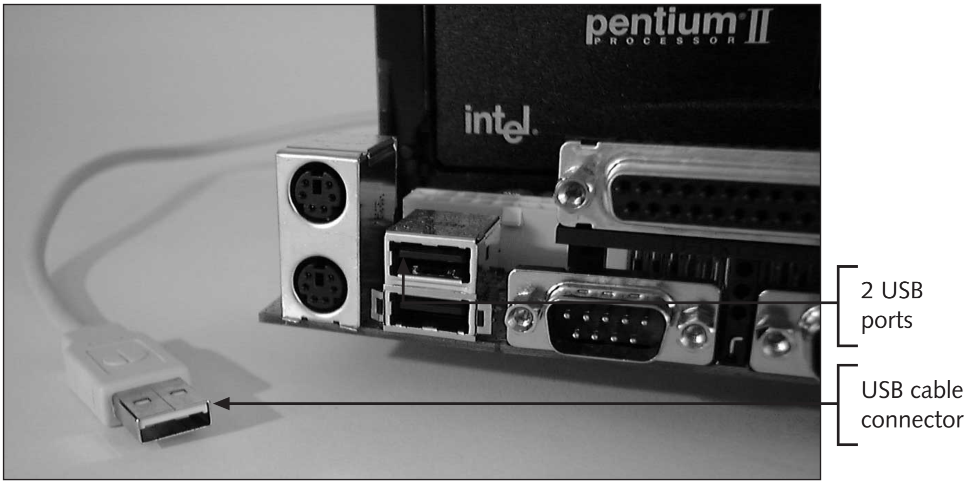

Figure 3-20 A system board with two USB ports and a USB cable; note the rectangular shape of the connection as compared to the nearby serial and parallel D-shaped ports

A **USB host controller**, which for the 400 series Intel chip set, is included in the PCI controller chip, manages the USB bus. As many as 127 USB devices can be daisy-chained together using USB cables up to five meters long. The host controller manages communication to the CPU for all devices, using only a single IRQ, I/O address range, and DMA channel. USB allows for **hot-swapping**, meaning that a device can be plugged into a USB port while the computer is running, and the host controller will sense the device and configure it without your having to reboot the computer. One USB device, such as a keyboard, can provide a port for another device, or a device can serve as a **hub**, allowing several devices to connect to it. There can also be a standalone hub into which several devices can be plugged. In USB technology, the host controller polls each device, asking if data is ready to be sent or requesting to send data to the device. The USB cable has four wires, two for power and two for communication. The two power wires (one carries voltage and the other is ground) allow the host controller to provide power to a device.

I/O devices that are now or are soon to be available with a USB connection are the mouse, joystick, keyboard, printer, scanner, monitor, modem, video camera, fax machine, and digital telephone.

USB must be supported by the operating system in order to work. Windows 95 with the USB update, Windows 98, and Windows 2000 support USB, but Windows NT does not. For more information about USB, see the forum web site at *www.usb.org*.

FireWire or i.Link or 1394

A+ CORE 1.4, 1.7

FireWire and **i.Link** are the common names for another peripheral bus officially named **IEEE 1394** (or sometimes simply called 1394) after the group that designed the bus. The Institute of Electrical and Electronics Engineers was primarily led by Apple Computer and

A+CORE 1.4, 1.7

Texas Instruments in the initial design. FireWire is similar in design to USB, using serial transmission of data, but faster. FireWire supports data speeds as high as 1.2 Gbps (gigabits per second), much faster than USB. It is a viable option for connecting network cards, camcorders, DVD, and other high-speed, high-volume devices. Whereas USB is looking to replace slow serial and parallel ports, FireWire is likely to replace SCSI as a solution for high-volume multi-media external devices such as digital camcorders and hard drives. SCSI is a very fast, but difficult to configure, peripheral bus that will be discussed in Chapter 6.

Devices are daisy-chained together and managed by a host controller using a single set of system resources (an IRQ, an I/O address range, and a DMA channel). The one host controller can support up to 63 FireWire devices. Just as with USB, FireWire must be supported by the operating system. Windows 2000, Windows NT, and Windows 98 all support FireWire.

The 1394 Trade Association has developed a new standard, **IEEE 1394.3**, designed for peer-to-peer data transmission. Using this new standard, imaging devices, such as scanners and digital cameras, can send images and photos directly to printers without involving a computer.

Local I/O Buses

A+CORE 4.3

Recall that the primary intent of a local bus is to provide direct access to the CPU for a few fast devices, such as memory and video, that can run at nearly the same speed as the CPU. A local I/O bus must connect to the CPU by way of the memory bus. Figure 3-21 shows an example of a proprietary local bus that uses a 32-bit expansion slot created by adding an extra slot to a 16-bit ISA slot.

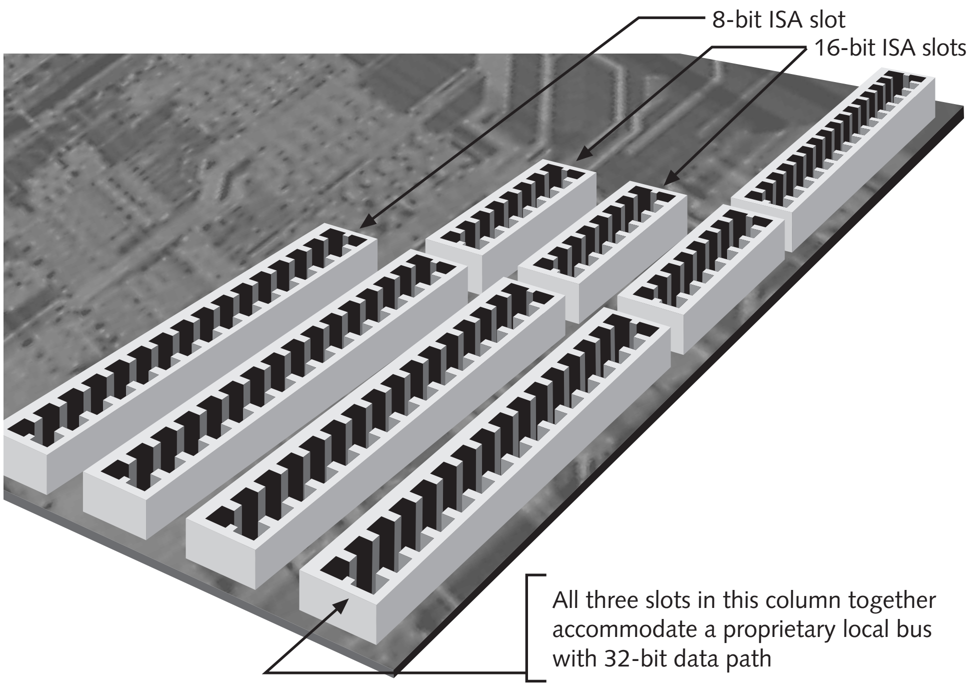

Figure 3-21 Three kinds of legacy bus connections on the same system board

A+ CORE 4.3

In an attempt to create a standard for local 32-bit buses, many manufacturers endorsed the **VESA (Video Electronics Standards Association)** VL bus. Many system boards offered the VESA local bus for video and memory circuit boards. The expansion slot for a VESA local bus includes the 16 bits for the ISA slot plus an added extension with another 116 pins (see Figure 3-22). The VESA bus has been replaced with the PCI bus.

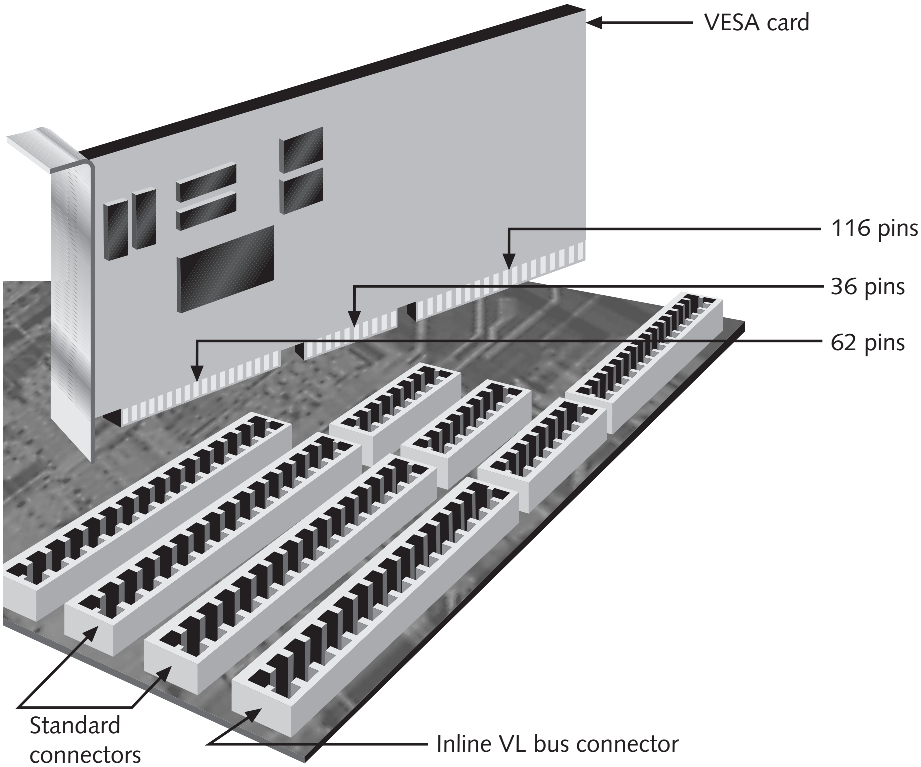

Figure 3-22 VESA local bus expansion slot

PCI Bus

Another local I/O bus, the **PCI local bus** (**peripheral component interconnect bus**) is now the standard local I/O bus not only with Pentium CPUs but also with RISC CPUs. Standard PCI has a 32-bit data path and runs at 33 MHz when the system board runs at 66 MHz. However, the PCI specifications can also use a 64-bit data path and can run at a speed of 66 MHz when the system bus runs at 133 MHz. Also, an addendum to the PCI specifications, called PCI-X, released in September, 1999, enables PCI to run at 133 MHz.

A+ CORE 4.3

One advantage of the PCI local bus is that devices connected to it can run at one speed while the CPU runs at a different speed. Devices connected to the VESA bus must run at the same clock speed as the CPU, which forces the CPU to endure frequent wait states. The PCI bus expansion slots are shorter than ISA slots (see Figure 3-23) and set a little farther away from the edge of the system board. Figure 3-24 shows the pinouts for the standard PCI slot.

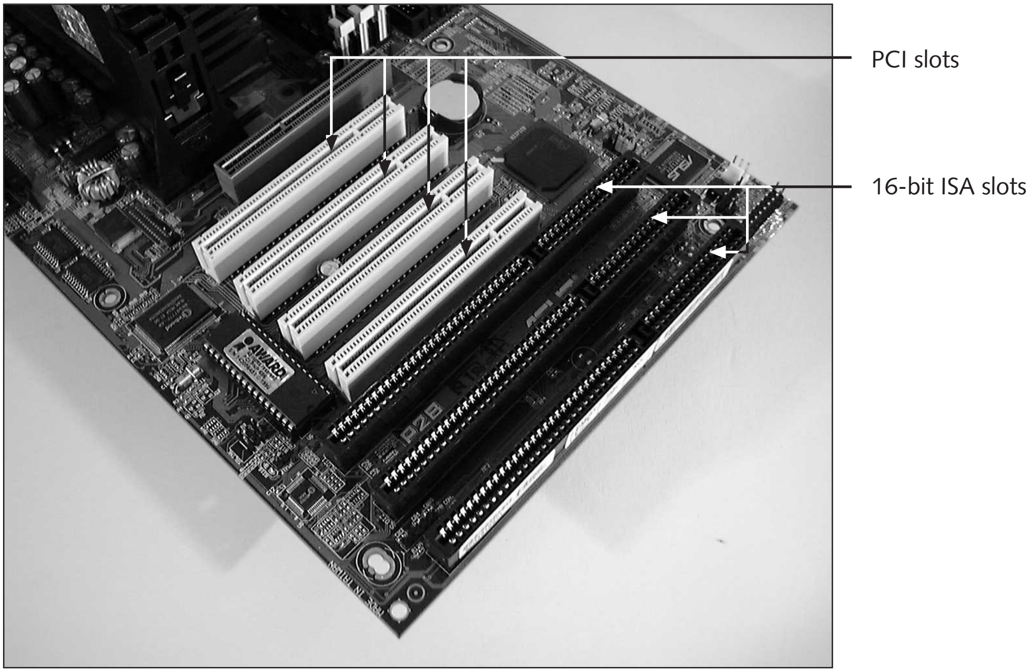

Figure 3-23 PCI bus expansion slots are shorter than ISA slots and offset farther

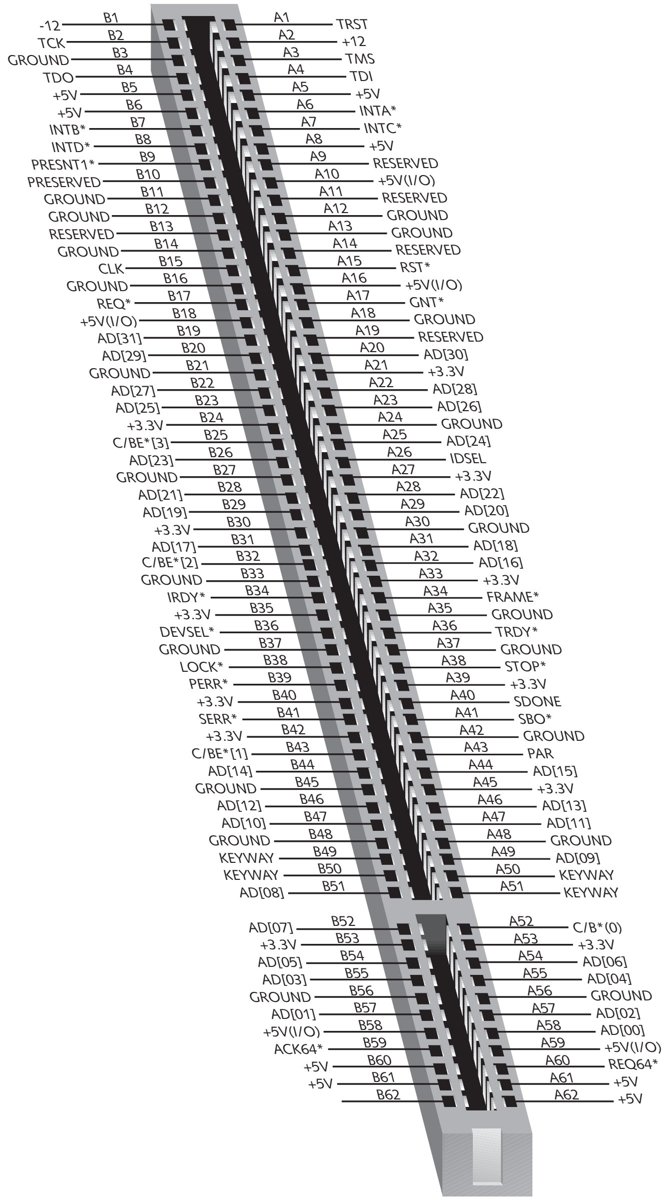

Figure 3-24 Standard PCI slot pinouts

As you learned earlier in the chapter, in addition to supporting I/O devices connected to it, the PCI bus also serves another function for the Intel 400 series chip set. The PCI bus interfaces with the expansion bus and the memory bus, serving as the go-between for the two, controlling the input and output to the expansion bus. As you can see in Figure 3-25, the memory bus is isolated from the ISA bus by the PCI bus. The connection between the

two is the PCI bridge. The bridge allows the PCI bus to control the traffic not only from its own local devices but also from the ISA bus.

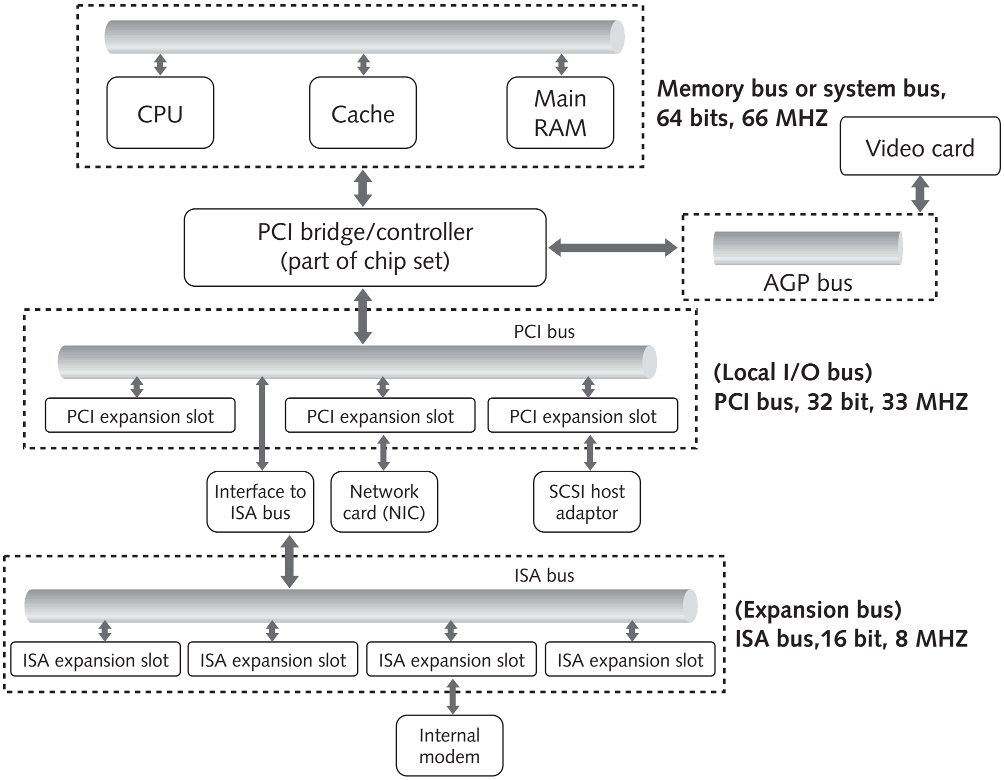

Figure 3-25 The PCI bus serves as the middleman between the memory bus and the expansion bus

In Figure 3-25, the SCSI (Small Computer System Interface) host adapter (discussed at length later in the book) and a network interface card (NIC) are connected to the PCI bus. (Physically, each card is inserted in a PCI expansion slot.) The PCI bridge/controller accesses the local bus where the CPU and memory are allowed to run at top speed without interference or wait states. If the CPU wants to send data to the network card, for example, it dumps it on the PCI bridge/controller at top speed. The controller puts the data in its own buffer or temporary memory storage and then writes it to the network card at a pace slower than the local bus. The bridge/controller eliminates interference with the local bus.

Figure 3-25 also shows the interface from the PCI bus to the ISA bus. This interface is a significant feature that distinguishes the PCI bus from other buses. The PCI bus was not designed to replace the traditional expansion bus, but to support it. The ISA bus in the diagram passes data through the interface to the PCI bus, which in turn passes the data on to the memory bus, to the CPU, and to memory.

The PCI bus also supports bus mastering. A bus master is an intelligent device (i.e., it has a microprocessor installed that manages the device) that, when attached to the PCI bus, can gain access to memory and other devices on the bus without interrupting the action of the CPU. The CPU and the bus mastering devices can run concurrently and independently of each other.

Because of the effective design of the PCI bus, the throughput performance, or the data transfer rate per second, is 132 MB when the bus is running at 33 MHz with a 32-bit data path. **Throughput performance**, or **data throughput**, is a measure of the actual data transmitted by the bus, not including error-checking bits or redundant data.

Accelerated Graphics Port

A+CORE 4.3

The **accelerated graphics port** (**AGP**) is designed to provide fast access to video. System boards have a single AGP slot to support an AGP video card (see Figure 3-26). AGP is more of a port than a bus, since it does not allow for expandability and can only support a single card. Notice in the diagram in Figure 3-25 that the faster AGP bus has a direct connection to the CPU without having to use the slower PCI bus.

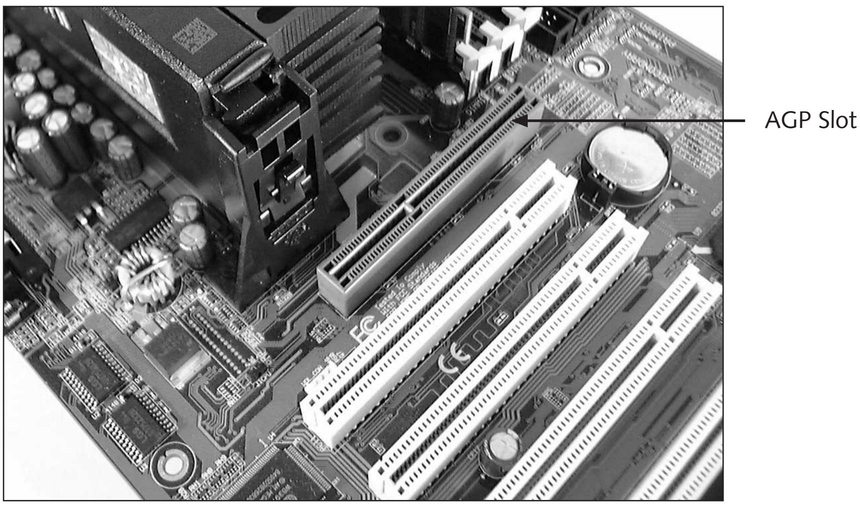

Figure 3-26 A system board will have only one AGP slot, which is used to support a video/graphics card

A+ CORE 4.3

TIP A slot or socket is the physical connector on a system board for a device and does not include logic to support the connected device. A port is a socket or slot but goes a step further and also includes the logic to support that connection (example: serial and parallel ports). A bus, among other things, provides the logic to support several devices connected to the system board, but does not include the physical connection itself. You need an expansion slot for that.

The AGP bus runs at the same speed as the memory bus, connects directly to it, and has a 32-bit-wide data path. AGP runs faster than PCI, running at half the memory bus speed, but also offers additional features that give it overall better performance for video than PCI. It offers an improved rendering of 3-D images when software is designed to use it.

AGP can share system memory with the CPU to do its calculations and, therefore, does not always have to first copy data from system memory to video memory on the graphics card. This feature, known as direct memory execute (DIME), is probably the most powerful feature of AGP. The first AGP specification defined AGP 2X, which allowed AGP to transfer two cycles of data during a single AGP clock beat. The AGP 2.0 specification defined AGP 4X whereby four cycles of data can be transferred during a single AGP clock beat yielding an overall data throughput of more than 1 GB/sec (gigabytes per second).

Figure 3-26 shows a 132-pin AGP slot on a system board. The latest AGP standard, called the AGP Pro, has provision for a longer slot. The new 188-pin slot has extensions on both ends that each contain an additional 28 pins that are used to provide extra voltage to the AGP video card in the slot. AGP Pro is used for high-end workstations that require powerful graphic accelerator cards for graphic-intensive applications.

In order for AGP to work at its full potential, the system board must be running at a minimum of 100 MHz, and the operating system must support AGP. Windows 98 and Windows 2000 both support AGP. See *www.agpforum.org* and *developer.intel.com/technology/agp/* for more information.

Audio Modem Riser

Newer system boards sometimes have an audio modem riser (AMR) slot that can accommodate a small modem card or sound card. These small cards are inexpensive since most of the logic to support audio or the modem is contained within the system board chip set. The AMR slot makes it possible to add the card at a low cost without using up a PCI or ISA slot.

Setting the CPU and Bus Speeds

A+CORE 1.8

You can, to some extent, control the speed of your system. Table 3-11 lists how the CPU and several bus speeds are controlled. There are two ways you can change the speed of a computer:

1. Change the speed of the memory bus. Whatever the memory bus speed is, the PCI bus speed is half or one third of that.
2. Change the multiplier that determines the speed of the CPU. The choices for the multiplier normally are 1.5, 2, 2.5, 3, 3.5 and so forth.

Table 3-11 System-board speeds and how they are determined

Bus or Device	How Speed Is Determined	How Controlled
CPU	Processor speed = memory bus speed x multiplier. Typical speeds are 350 MHz, 450 MHz, and 500 MHz	Multiplier is set by jumpers or DIP switches on the system board or in CMOS setup
Memory bus or system bus	System board manufacturer recommends the speed based on the processor and processors rated speed. Typical values are 66 MHz, 100 MHz, and 133 MHz	Set by jumpers, DIP switches, or in CMOS setup. Most commonly set by jumpers.
PCI bus	Memory bus speed / 2 (or for faster boards, it can be divided by 3)	The speed is set when you set the speed of the memory bus; either 33 MHz or 66 MHz
ISA bus	Runs at only one speed: 8.77 MHz	NA

Studies have shown that when the multiplier is large, the overall performance of the system is not as good as when the multiplier is small. This is a reasonable result because you are interested in the overall speed of the computer, which includes the CPU and the buses, not just the speed of the CPU. For example, a bus speed of 60 MHz and a multiplier of 5 yield a relatively fast CPU but a relatively slow bus. It's better to have a bus speed of 80 MHz and a multiplier of 3 so that the bus is running fast enough to keep up with the CPU.

See the system board documentation to learn how to set these speeds using jumpers, DIP switches, or CMOS setup. Figure 3-27 shows the documentation for one system board that uses one bank of jumpers to set the CPU-to-bus multiplier and another jumper bank to set the bus frequency. The steps to do that are:

1. Read the documentation of your CPU to determine its recommended frequency. In our example, we are using a Pentium II rated for 350MHz. Find the row labeled Pentium 350MHz in the table of Figure 3-27.

2. Read the multiplier from the selected row, which is 3.5x. Find the jumper settings for a multiplier of 3.5 in the possible jumper combinations for the CPU Core: BUS Frequency Multiple (fourth entry in first row).
3. The CPU type and speed also determine the bus frequency, which is 100MHz. To set the bus frequency to 100 MHz, find the jumper combination for the second jumper bank, which is the fourth entry in the list of selections for the CPU External Clock (Bus) Frequency Selection.
4. Set the jumpers in the two jumper banks. Figure 3-28 shows the jumper group for the multiplier set to 3.5.

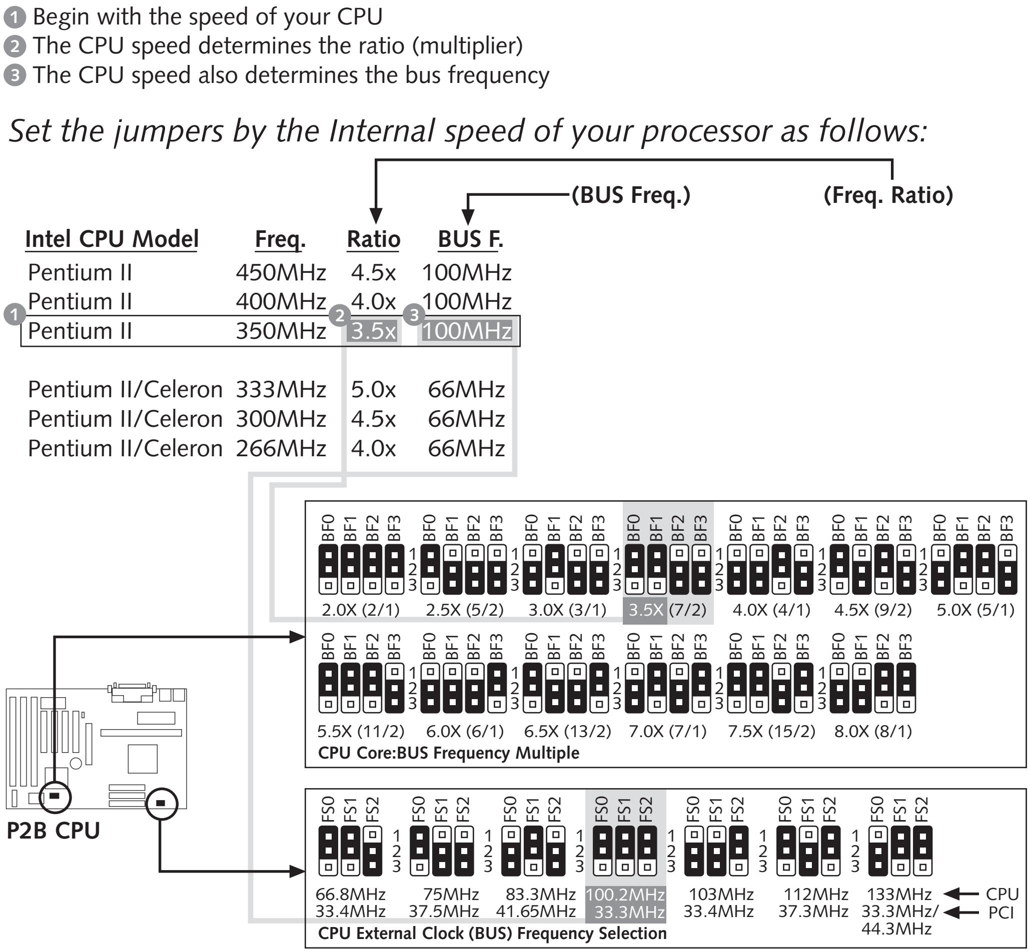

Figure 3-27 Based on the advertised speed of your CPU, select the multiplier and the bus frequency from the table, which then determines the jumper settings to use

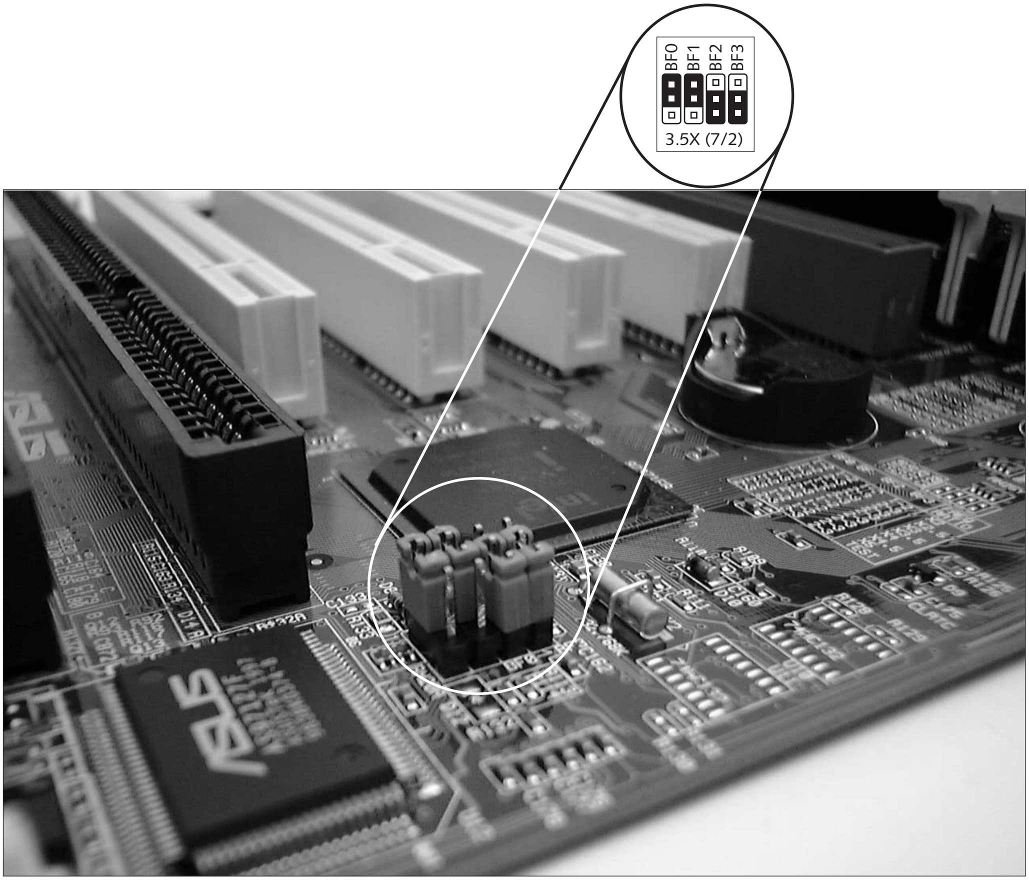

Figure 3-28 Jumper group that controls the CPU core-to-bus frequency. Compare this photo to the diagram in Figure 3-27. The jumpers are set for the multiplier = 3.5.

On-Board Ports

Many system boards contain **on-board ports** such as a keyboard port and a mouse port. In addition, a parallel printer port and one or two serial ports might be located directly on the system board. Few older system boards contain more ports than these. Some systems also have a video or network port, and newer system boards contain one or two USB ports.

You don't have to replace an entire system board if one port fails. Most system boards contain jumpers or DIP switches that can tell the CPU to disable one port and look to an expansion card for the port instead. Ports can also be disabled through CMOS setup.

When buying a new computer or system board, look for the ability to disable ports, floppy drive connectors, or hard drive connectors coming directly from the system board by changing the hardware configuration. You can easily tell if ports on the outside of the case

are directly connected to the system board without opening the case; the ports are lined up along the bottom of the computer case, as shown in Figure 3-29.

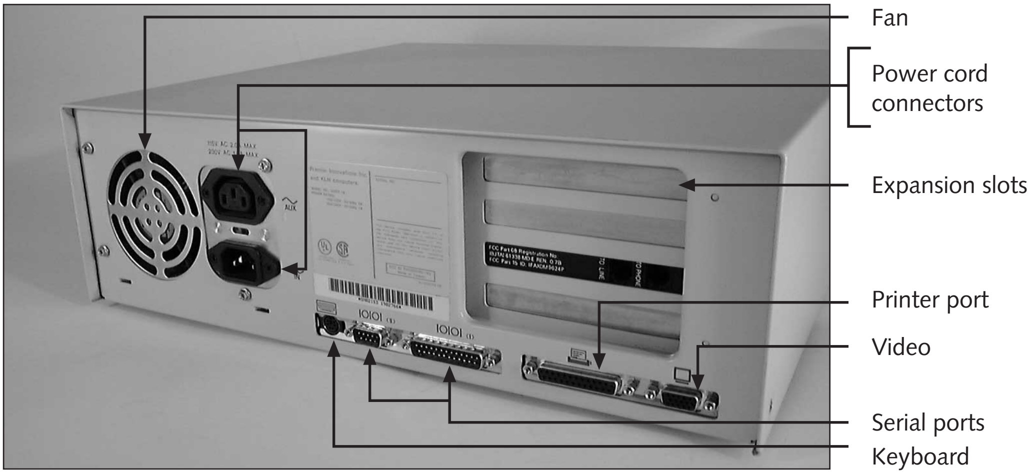

Figure 3-29 Ports along the bottom of the computer case usually come directly off the system board

HARDWARE CONFIGURATION

Recall that hardware configuration information communicates to the CPU what hardware components are present in the system and how they are set up to interface with the CPU. Hardware configuration includes information such as how much memory is available, what power management features are present, and whether disk drives, hard drives, modems, serial ports, and the like are connected. Remember that during POST, BIOS looks to the system configuration information to determine the equipment it should expect to find and how that equipment interfaces with the CPU. The CPU uses this information later to process data and instructions. Configuration information is provided on the system board in three different ways, as discussed in Chapter 2: DIP switches, jumpers, and CMOS. Earlier in the chapter you saw several examples of how jumpers are used to set configuration information on the system board. Here we look more closely at the CMOS setup screens.

Setup Stored on a CMOS Chip

A+CORE 1.1, 4.4

Computers today store most configuration information on one CMOS chip that retains the data even when the computer is turned off. (There are actually many CMOS chips on a system board, used for various purposes.) A battery near the CMOS chip provides enough electricity to enable the chip to maintain its data. If the battery is disconnected or fails, setup information is lost. Password information is also a part of the computer's setup that is stored in CMOS. The program to change the setup information is now stored in ROM but once was on a disk that came with the computer.

A+ CORE 1.1, 4.4

System-board manuals should contain a list of all CMOS settings, an explanation of their meanings, and their recommended values. When you purchase a system board or a computer, be sure the manual is included for this purpose. If you don't have the manual, you can sometimes go to the system-board manufacturer's Web site and download the information you need to understand the specific CMOS settings of your computer.

Some CMOS settings are listed in Table 3-12.

Table 3-12 CMOS settings and their purpose

Category	Setting	Description
Standard CMOS Setup	Date and time	Use to set system date and time (called the real time clock).
	Primary display	Use to tell POST and DOS (but not Windows) the type of video being used.
	Keyboard	Use to tell system if keyboard is installed or not installed. Useful if the computer is used as a print or file server and you don't want someone changing settings.
	Hard disk type	Use to record size and mapping of the drive.
	Floppy disk type	Choices are usually 3½ inch and 5¼ inch.
Advanced CMOS Setup	Above 1 MB memory test	Use to disable POST check of this memory to speed up booting. The OS will check this memory anyway.
	Memory parity error check	If you have a parity system board, use to enable parity checking to ensure that memory is correct.
	Numeric processor test	Enabled unless you have an old 386 or 486SX computer
	System boot sequence	Use to establish the drive the system turns to first to look for an OS. Normally drive A, then C.
	External cache memory	Use to enable if you have L2 cache. A frequent error in setup is to have cache but not use it because it's disabled here.
	Internal cache memory	Normally enabled; disable only for old 386 computers.
	Password checking option	Use to establish a startup password. Use this only if you really have a problem with someone using your PC who can't be trusted.
	Video ROM shadow C000, 16K	For DOS and Windows 9x, shadowing video ROM is recommended because ROM runs slower than RAM.
	System ROM shadow F000, 64K	Enabling shadow system ROM is recommended.
	IDE Multi-block mode	Enables a hard drive to read or write several sectors at a time. Dependent on the kind of hard drive you have.
	Boot sector virus protection	Gives a warning when something is being written to the boot sector of the hard drive. Can be a nuisance if your software is designed to write to the boot sector regularly.

Table 3-12 CMOS settings and their purpose (continued)

Category	Setting	Description
Advanced Chip Set Setup	AT bus clock selection	Gives the number by which the CPU speed is divided to get the ISA or EISA bus speed
	ISA bus speed	Gives the number by which the PCI bus speed is divided to get the ISA bus speed
	Bus mode	Can be set to synchronous or asynchronous modes. In synchronous mode, the bus uses the CPU clock. In asynchronous mode its own AT bus clock is used.
	AT cycle wait state	The number of wait states the CPU must endure while it interfaces with a device on the ISA or EISA bus. Increase this if an old and slow ISA card is not working well.
	Memory read wait state	Number of wait states the CPU must endure while reading from RAM
	Memory write wait state	Number of wait states the CPU must endure while writing to RAM
	Cache read option	Sometimes called "cache read hit burst"
		The number of clock beats needed to load four 32-bit words into the CPU's internal cache. 4-1-1-1 is the usual choice.
	Fast cache read/write	Refers to external cache. Enable it if you have two banks of cache, 64K or 256K.
	Cache wait state	Refers to external cache. The number of wait states the CPU must use while accessing cache.
Power Menu	Power Management	Disable or enable all power management features. These features are designed to conserve electricity.
	HDD Power Down	Disable or enable the feature to shut down the hard drive after a period of inactivity
	Wake on LAN	Wake on LAN allows your PC to be booted from another computer on the same network. It requires an ATX power supply that supports the feature.
	Wake on keyboard	Allows you to power up your PC by pressing a certain key combination.

Chapter Summary

- The system board is the most complicated of all the components inside the computer. It contains the CPU and accompanying chip set, the real-time clock, ROM BIOS, CMOS configuration chip, RAM, RAM cache, system bus, expansion slots, jumpers, ports, and power supply connections. The system board you select determines both the capabilities and limitations of your system.
- The most important component on the system board is the CPU, or central processing unit. The CPU is the microprocessor at the heart of a PC system, where almost all operations must ultimately be processed. The CPU is rated according to its speed,

efficiency of programming code, word size, data path size, maximum memory addresses, size of internal cache, multiprocessing abilities, and special functions. Earlier Intel CPUs include the 80386DX, 80386SX, 80486DX, and 80486SX. The latest family of Intel CPUs is the Pentium family, including the Classic Pentium, Pentium MMX, Pentium Pro, Pentium II, Celeron, and Pentium III. AMD and Cyrix are Intel's chief competitors for the CPU market. CPUs can use either RISC or CISC technology or a combination of the two.

- Newer CPUs require extra cooling, which can be accomplished with a CPU heat sink and cooling fan located on top of or near the CPU.
- The common CPU sockets and slots today are Socket 7, Socket 370, Slot 1, Slot A, and Slot 2. A slot looks like an expansion slot.
- Because some CPUs require one voltage for internal core operations and another voltage for external I/O operations, system boards might have a voltage regulator on board.
- Some components can be built into the system board, in which case they are called on-board components, or they can be attached to the system in some other way such as on an expansion card.
- ROM chips contain the programming code to manage POST and system BIOS and to change the CMOS settings. The setup or CMOS chip holds configuration information.
- A chip set is a group of chips on the system board that supports the CPU. Intel is the most popular manufacturer of chip sets.
- The total BIOS of a system includes the ROM BIOS on the system board as well as BIOS on expansion cards. Plug and Play BIOS is designed to work in harmony with Windows 9x or Windows 2000 to resolve resource conflicts from expansion cards and other devices. Flash ROM allows the ROM BIOS to be upgraded without having to change the ROM chip.
- Dynamic RAM (DRAM) is slower than static RAM (SRAM) because dynamic RAM must be refreshed. RAM usually comes packaged as a SIMM or DIMM memory module.
- Two kinds of static RAM cache for the slower DRAM are internal and external cache, sometimes called Level 1 and Level 2 cache.
- Level 1 cache is contained on the CPU microchip, and level 2 cache is external to this microchip.
- A bus is a path on the system board that carries electrical power, control signals, memory addresses, and data to different components on the board.
- A bus can be 16, 32, 64 or more bits wide. The first ISA bus had an 8-bit data path. The second ISA bus had a 16-bit data path
- Some well-known buses are the 16-bit ISA bus, 32-bit MCA and EISA buses, and the two local buses, the VESA bus and the PCI bus. A local bus is designed to allow fast devices quicker and more direct access to the CPU than that allowed by other buses.
- The VESA local bus is a standard designed by the Video Electronics Standards Association. The PCI bus is presently the most popular local bus. To gain the maximum overall computer performance, the multiplier relating the bus speed to the CPU speed should be small.

- Expansion slots can be located on the system board, but they are sometimes stacked vertically in the computer case on a second board devoted to that purpose.
- Jumpers on the system board can be used to set the system-board speed and the CPU multiplier that determines the CPU speed.
- Sometimes the CPU must be slowed down to accommodate slower devices, enduring wait states that cause it to wait one clock beat. Wait states often mean a significant reduction in performance.

KEY TERMS

Accelerated graphics port (AGP) — A slot on a system board for a video card that provides transfer of video data from the CPU that is synchronized with the memory bus.

Advanced Transfer Cache (ATC) — A type of L2 cache contained within the Pentium processor housing that is embedded on the same core processor die as the CPU itself.

Backside bus — The bus between the CPU and the L2 cache inside the CPU housing.

Bus — Strips of parallel wires or printed circuits used to transmit electronic signals on the system board to other devices. Most Pentium systems use a 32-bit or 64-bit bus.

Bus speed — The speed or frequency at which the data on the system board is moving.

Chip set — A set of chips on the system board that collectively controls the memory cache, external buses, and some peripherals.

Clock speed — The speed or frequency that determines the speed at which devices on the system bus operate, usually expressed in MHz. Different components on a system board operate at different speeds, which are determined by multiplying or dividing a factor by the clock speed. The clock speed is itself determined by a crystal or oscillator located somewhere on the system board.

Data path — The number of bits of data transmitted simultaneously on a bus. The size of a bus, such as a 32-bit-wide data path in a PCI bus.

DIMM (dual in-line memory module) — A miniature circuit board that holds memory chips and has a 64-bit data path. Because Pentium system boards also use a 64-bit memory bus, it is possible to use only a single DIMM on these system boards.

Discrete L2 cache — A type of L2 cache contained within the Pentium processor housing, but on a different die, with a cache bus between the processor and the cache.

Dual voltage CPU — A CPU that requires two different voltages, one for internal processing and the other for I/O processing.

Dynamic RAM (DRAM) — The most commonly used type of system memory, it requires refreshing every few milliseconds.

EISA (Extended Industry Standard Architecture) bus — A 32-bit bus that can transfer 4 bytes at a time at a speed of about 20 MHz.

ESCD (extended system configuration data) — A list written to the BIOS chip of information about legacy devices that Plug and Play uses to configure these devices.

Expansion bus — A bus that does not run synchronized with the system clock.

External cache — Static cache memory, stored on the system board or inside CPU housing, that is not part of the CPU (also called level 2 or L2 cache).

Field replaceable unit — A component in a computer or device that can be replaced with a new component without sending the computer or device back to the manufacturer. Example: a DIMM memory module on a system board.

FireWire — An expansion bus that can also be configured to work as a local bus. It is expected to replace the SCSI bus, providing an easy method to install and configure fast I/O devices. Also called IEEE 1394 and i.Link.

Frontside bus — The bus between the CPU and the memory outside the CPU housing.

Heat sink — A piece of metal, with cooling fins, that can be attached to or mounted on an integrated chip (such as the CPU) to dissipate heat.

Hot-swapping — When a device can be plugged into a computer while it is turned on and the computer will sense the device and configure it without rebooting.

Instruction set — The set of instructions, on the CPU chip, that the computer can perform directly (such as ADD and MOVE).

Internal cache — Memory cache that is faster than external cache, and is contained inside 80486 and Pentium chips (also referred to as primary, Level 1, or L1 cache).

IEEE 1394 — *See* Fire Wire.

ISA bus — An 8-bit industry standard architecture bus used on the original 8088 PC. Sixteen-bit ISA buses were designed for the 286 AT, and are still used in Pentiums for devices such as modems.

Level 1 cache — *See* Internal cache.

Level 2 cache — *See* External cache.

Local bus — A bus that operates at a speed synchronized with the CPU speed.

Local I/O bus — A local bus that provides I/O devices with fast access to the CPU.

Low insertion force (LIF) — A socket feature that requires the installer to manually apply an even force over the microchip when inserting the chip into the socket.

MCA (Micro Channel Architecture) bus — A proprietary IBM PS/2 bus, seldom seen today, with a width of 16 or 32 bits and multiple master control, which allowed for multitasking.

Memory bus — The bus between the CPU and memory on the system board. Also called the system bus or the host bus.

Memory cache — A small amount of faster RAM that stores recently retrieved data, in anticipation of what the CPU will request next, thus speeding up access.

Multiplier — The factor by which the bus speed or frequency is multiplied to get the CPU clock speed.

North bridge — That portion of the chip set hub that connects faster I/O buses (for example, AGP bus) to the system bus. Compare to South bridge.

On-board ports — Ports that are directly on the system board, such as a built-in keyboard port or on-board serial port.

Overclocking — Running a system board at a speed that is not recommended or guaranteed by CPU or chip set manufacturers.

P1 connector — Power connection on an ATX system board.

Parity — An error-checking scheme in which a ninth, or "parity," bit is added. The value of the parity bit is set to either 0 or 1 to provide an even number of ones for even parity and an odd number of ones for odd parity.

Parity error — An error that occurs when the number of 1s in the byte is not in agreement with the expected number.

PCI (peripheral component interconnect) bus — A bus common on Pentium computers that runs at speeds of up to 33 MHz or 66 MHz, with a 32-bit-wide or 64-bit-wide data path. For most chip sets, it serves as the middle layer between the memory bus and expansion buses.

Pin grid array (PGA) — A feature of a CPU socket whereby the pins are aligned in uniform rows around the socket.

Plug and Play — A technology in which the operating system and BIOS are designed to automatically configure new hardware devices to eliminate system resource conflicts (such as IRQ and port conflicts).

Primary cache — *See* Internal cache.

Processor speed — The speed or frequency at which the CPU operates. Usually expressed in MHz.

Refresh — The process of periodically rewriting the data for instance, on dynamic RAM.

RISC (reduced instruction set computer) chips — Chips that incorporate only the most frequently used instructions, so that the computer operates faster (for example, the PowerPC uses RISC chips).

SECC (Single Edge Contact Cartridge) — A type of cartridge that houses the Pentium III processor.

SC330 (Slot Connector 330) — A 330-pin system board connector used to contain the Pentium III Xeon. Also called Slot 2.

SIMM (single in-line memory module) — A miniature circuit board that holds memory chips and has a 32-bit data path. Because Pentium system boards use a 64-bit memory bus, you must install two SIMMs at a time.

Single voltage CPU — A CPU that requires one voltage for both internal and I/O operations.

South bridge — That portion of the chip set hub that connects slower I/O buses (for example, ISA bus) to the system bus. Compare to North bridge.

Staggered pin grid array (SPGA) — A feature of a CPU socket whereby the pins are staggered over the socket in order to squeeze more pins into a small space.

Static RAM (SRAM) — RAM chips that retain information without the need for refreshing, as long as the computer's power is on. They are more expensive than traditional DRAM.

System bus — Today the system bus usually means the memory bus. However, sometimes it is used to refer to other buses on the system board. *See* memory bus.

Turbo mode — A means of doubling the external clock speed by pressing a button on the case of some older computers.

Universal serial bus (USB) — A bus that is expected to eventually replace serial and parallel ports, designed to make installation and configuration of I/O devices easy, providing room for as many as 127 devices daisy-chained together. The USB uses only a single set of resources for all devices on the bus.

VESA (Video Electronics Standards Association) VL bus — An outdated local bus used on 80486 computers for connecting 32-bit adapters directly to the local processor bus.

Wait state — A clock tick in which nothing happens, used to ensure that the microprocessor isn't getting ahead of slower components. A 0-wait state is preferable to a 1-wait state. Too many wait states can slow a system down.

Zero insertion force (ZIF) — A socket feature that uses a small lever to apply even force when installing the microchip into the socket.

Review Questions

1. What are the two most popular types of system-board form factors?
2. How many power cords connect to a Baby AT system board have?
3. Name 10 components that are contained on a system board.
4. When is it appropriate to have a Slot 1 on a system board?
5. Why would you want both ISA and PCI expansion slots on a system board?
6. When people speak of bus size, what are they specifically referring to?
7. What characteristics of the system-board architecture determine the amount of memory that a CPU can address?
8. What was the first Intel CPU to contain internal cache?
9. If you know the system bus speed, how can you determine the CPU speed?
10. When is it appropriate to use a Celeron rather than a Pentium II in a computer system?
11. Which is more powerful, the Celeron or the Xeon processor?
12. What are the two major competitors of Intel in the CPU market?
13. Why did the competitors of the Intel Pentium II choose to stay with Socket 7 rather than use Slot 1 for their competing processors?
14. Contrast CISC and RISC technology.
15. In order to keep a CPU cool, which is better to use, and why: a heat sink or a cooling fan?
16. Describe the difference between a PGA socket and an SPGA socket.
17. Name a CPU that requires dual voltage. How are the two voltages used?
18. Why is AGP technology described as being more like a port than a bus?
19. What are the speeds of the most popular system boards currently available on the market?
20. Name three manufacturers of system-board chip sets.
21. Name the three most popular manufacturers of system BIOS programs.
22. When is Plug and Play not helpful in resolving a resource conflict?
23. Explain the difference between a local bus and an expansion bus.
24. How many bits are needed to store one byte when parity is used?
25. Why does memory need to be refreshed?

26. What is the name for the bus that connects L2 cache to the CPU inside the Pentium II processor housing?
27. Why don't all buses on a system board operate at the same speed?
28. Which is faster, the memory bus or the ISA bus?
29. What are the four categories of cargo that are carried over a bus?
30. Draw and label the ports that you find on the back of a typical computer.
31. Look in CMOS setup on your PC and list five things that you can change.
32. How can bus mastering improve overall computer performance?
33. What advantage does a USB have over serial ports?

PROJECTS

CAUTION

Unless you follow proper procedures, working inside your computer can cause serious damage—to both you and your computer. To ensure safety in your work setting, follow every precaution listed in the *Read This Before You Begin* section following this book's Introduction.

Hands-on Project

A+ CORE 3.2

Important Safety Precautions

In some of the following activities, you remove the cover of a computer and examine the components. Before you perform the exercises in this chapter, carefully read the following precautions and procedures to protect yourself and the equipment. Remember that you can compound a problem, causing even more damage, by carelessly neglecting these safety precautions.

The most common threat to hardware is electrostatic discharge (ESD), commonly known as static electricity. Damage by ESD can cause a catastrophic failure, which can destroy components, or can cause an upset failure that produces unpredictable malfunctions of components, which are often difficult to detect or diagnose.

The three best protections against ESD as you work on a computer are a ground strap, a ground mat, and static shielding bags. A ground bracelet, sometimes called a ground strap or a static strap (see Figure 3-30), is worn on your wrist and is grounded to a ground mat, computer case, or a ground prong of a wall outlet. It contains a current-eliminating device, called a resister, that prevents electricity from flowing through the bracelet to you. A ground mat (see Figure 3-31) often comes equipped with a cord to plug into the ground prong of the wall outlet and a snap on the mat to which you can attach the end of your ground strap. New components come shipped in static shielding bags. Save the bags to store other devices not currently installed on your PC.

A+ CORE 3.2

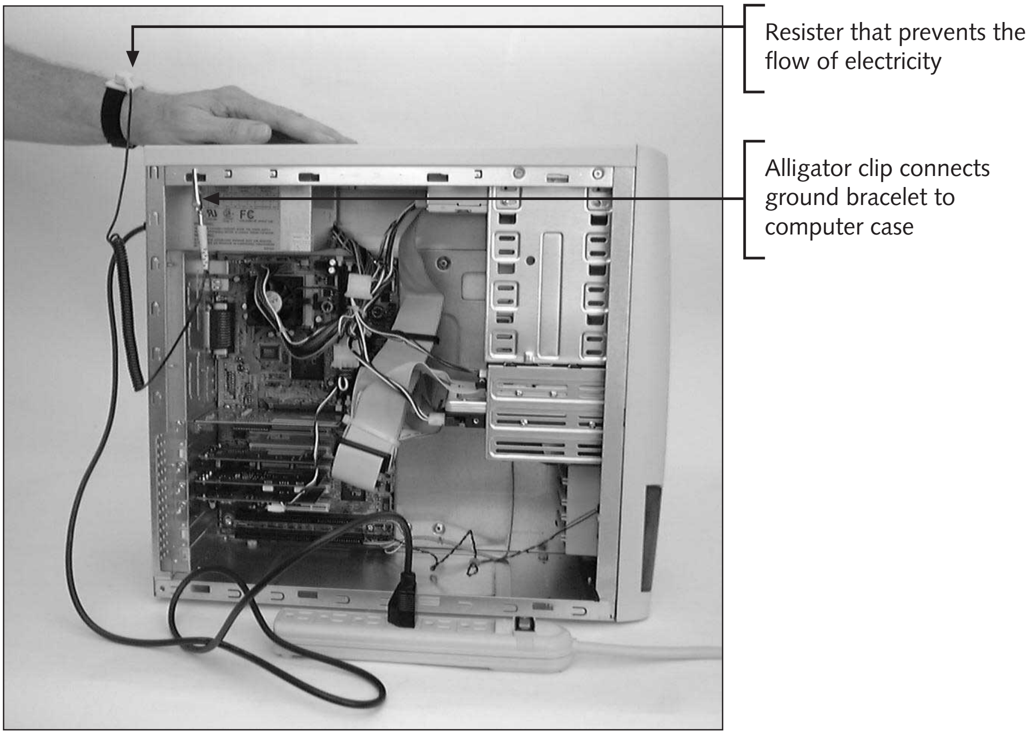

Figure 3-30 A ground bracelet, which protects against ESD, can clip to the side of the computer case and eliminates ESD between you and the case

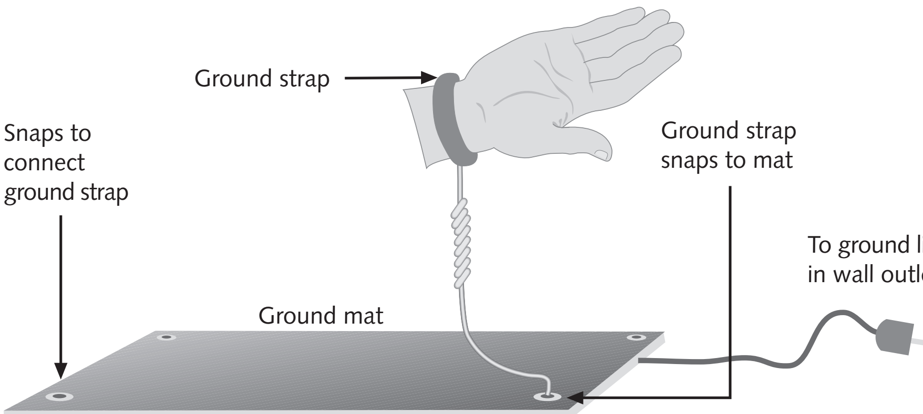

Figure 3-31 A ground bracelet can be connected to a ground mat, which is grounded by the wall outlet

Protect Against Electricity

A+CORE 3.2

You can actually destroy a computer chip with static electricity when you touch it, even though you might not feel a thing. You will learn more about this in Chapter 11, but for now, follow these rules to protect chips while you handle them:

1. Never touch the inside of a computer while it is turned on.
2. Never touch any component inside the computer without first grounding yourself to discharge any static electricity on your body. The best way to do this is to wear a ground bracelet or ground strap. If you don't have a ground bracelet, touch the metal case or power supply each time, before you touch any component, to discharge the electricity on your body.
3. Don't work on carpet. Work on a bare floor because a carpet collects static electricity, especially in cold weather.
4. Consider both the power supply and monitor to be a "black box." Don't open either unless you are trained to understand the dangers of and safety precautions for these devices. The power supply and the monitor have enough power inside them to kill you, even when they are unplugged.

Other Valuable Rules

1. When taking boards out of a PC, don't stack them. Stacking can loosen components.
2. Keep screws and spacers in an orderly place, such as a cup or tray.
3. Make notes as you work, so later you can backtrack.
4. In a classroom environment, after you have reassembled everything, before putting the cover back on, have your instructor check your work, *before* you power up.
5. Don't touch chips or edge connectors on boards unless absolutely necessary.
6. Don't touch chips with a magnetized screwdriver.
7. Don't use a graphite pencil to change DIP switch settings.
8. Don't put cards on top of or next to the monitor.
9. When setting down components, lay them on a grounded mat or static shielding bag.
10. Always turn off the PC before moving it (even a few inches) to protect the hard drive.
11. If you have been trained to work inside a monitor or power supply, while you are working inside either, be careful *not* to ground yourself.
12. When unpacking hardware or software, remove the packing tape from the work area as soon as possible.
13. Don't place a PC on the floor where it can be kicked.
14. Keep disks away from magnetic fields, heat, and extreme cold.
15. Don't open a disk's shuttle window or touch the surface of a disk.
16. Using a circuit tester, always verify that the ground plug in an outlet is physically grounded.

Examine the System Board

1. Look at the back of your computer. Without opening the case, list the ports that you believe to be coming directly from the system board.
2. Now look inside the case to verify your list.
   a. Follow these directions to remove the cover:
      - ❐ Turn off the PC and unplug it.
      - ❐ Unplug the monitor, mouse, and keyboard, and move them out of your way.
      - ❐ For a desktop case or tower case, locate and remove the screws on the back of the case. Look for the screws in each corner and one in the top center, as in Figure 3-28. Be careful that you don't unscrew any other screws besides these. The other screws probably are holding the power supply in place (see Figure 3-32).

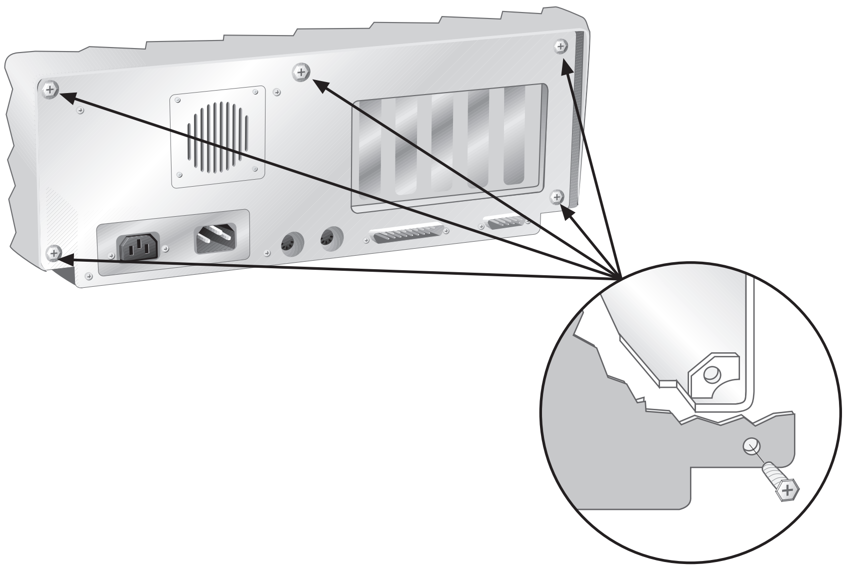

Figure 3-32 Locate the screws that hold the cover in place

- ❐ After you remove the cover screws, slide the cover forward and up to remove it from the case, as shown in Figure 3-33.

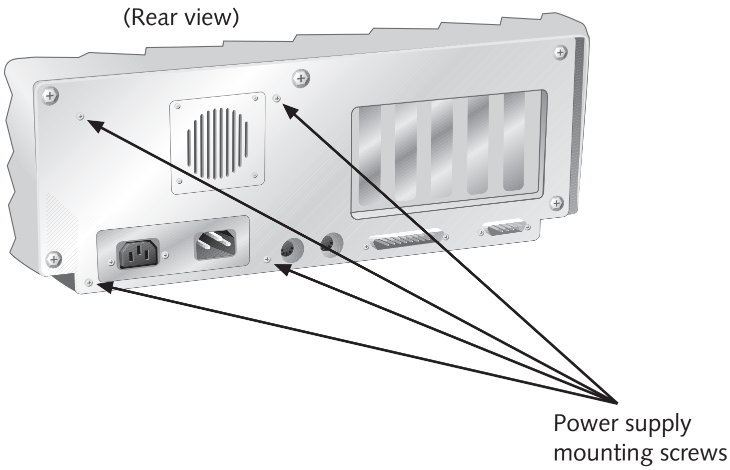

Figure 3-33 Power supply mounting screws

- For tower cases, the screws are also on the back. Look for screws in all four corners and down the sides (see Figure 3-34). Remove the screws and then slide the cover back slightly before lifting it up to remove it. Some tower cases have panels on either side of the case held in place with screws on the back of the case. Remove the screws and slide each panel toward the rear and then lift it off the case.

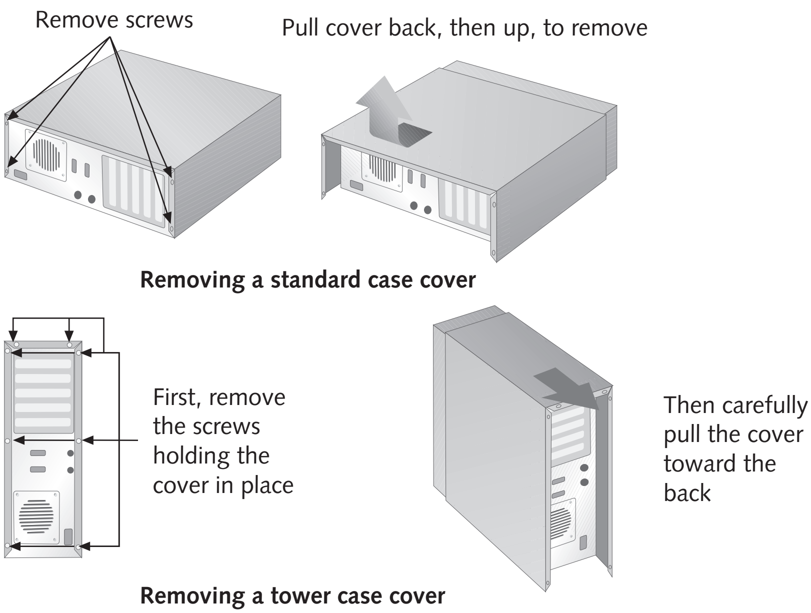

Figure 3-34 Removing the cover

b. Identify the following major components. Drawings in this and previous chapters should help.
   - ❒ Power supply
   - ❒ Floppy disk drive
   - ❒ Hard drive
   - ❒ System board

List the different circuit boards in the expansion slots. Was your guess correct about which ports come from the system board?

A+CORE 1.4

3. To expose the system board so you can identify its parts, remove all the expansion boards, following these procedures. (If you are working with a tower case, you can lay it on its side so the system board is on the bottom.)
   a. To make reassembling easier, take notes or make a sketch of the current placement of boards and cables. You can mark a cable on a card with a marker if you like. Note the orientation of the cable on the card. Each cable for the floppy disk drive, hard drive, or CD-ROM drive has a color on one side of the cable called the edge color. This color marks pin 1 of the cable. On the board, pin 1 is marked either as the number 1 or 2 beside the pin or, on the back side of the board, with a square soldering pad (see Figure 3-35).
   b. Remove the cables from the card. There is no need to remove the other end of the cable from its component (floppy disk drive, hard drive, or CD-ROM drive). Lay the cable over the top of the component or case.
   c. Remove the screw holding the board to the case.
   d. If you aren't wearing a ground bracelet, touch the case before you touch the board.
   e. Grasp the board with both hands and remove the board by lifting straight up, and rocking the board from end to end (not side to side). Rocking the board from side to side might spread the slot opening and weaken the connection.
4. Examine the board connector for the cable. Can you identify pin 1? Lay the board aside on a flat surface.
5. You probably will be able to see most if not all the components on the system board now without removing anything else. Draw a diagram of the system board and label these parts:
   - ❒ The CPU (include the prominent label on the CPU housing)
   - ❒ RAM (probably SIMMs or DIMMs)
   - ❒ Cache memory (probably one or more smaller SIMMs or DIMMs)
   - ❒ Expansion slots (identify the slots as ISA, EISA, MCA, PCI, VLB, etc.)
   - ❒ Each port coming directly from the system board
   - ❒ Power supply connections
   - ❒ ROM BIOS chip (Copy the writing on the top of the chip to paper. Identify the manufacturer, serial number, and date of manufacture of the chip.)

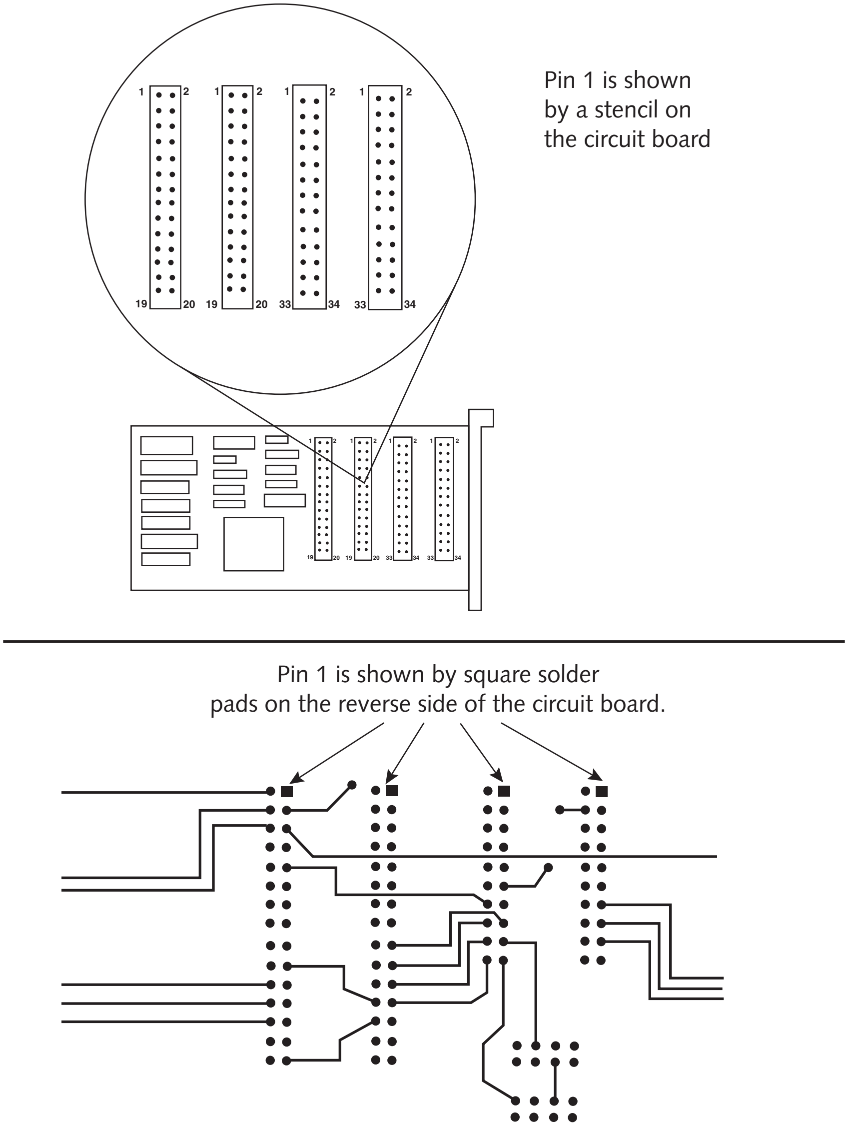

Figure 3-35 How to find pin 1 on an expansion card

6. Draw a rectangle on the diagram to represent each bank of jumpers on the board.
7. You can complete the following activity only if you have the documentation for the system board: locate the jumper or jumpers on the board that erases CMOS and/or the startup password, and label it on your diagram. It is often found near the battery. Some boards might not have one.

8. You are now ready to reassemble. Reverse the disassembling activities above. Place each card in its slot (it doesn't have to be the same slot, just the same bus) and replace the screw. Don't place the video card near the power supply.
9. Replace the cables, being sure to align the colored edge with pin 1. (In some cases it might work better to connect the cable to the card before you put the card in the expansion slot.)
10. Plug in the keyboard, monitor, and mouse.
11. In a classroom environment, have the instructor check your work before you power up.
12. Turn on the power and check that the PC is working properly before you replace the cover. Don't touch the inside of the case while the power is on.
13. If all is well, turn off the PC and replace the cover and its screws. If the PC does not work, don't panic! Just turn off the power and go back and check each cable connection and each expansion card. You probably have not solidly seated a card in the slot. After you have double-checked, try again.

Saving and Restoring CMOS Settings

Hands-on Project

In Chapter 2, you used Nuts & Bolts to record CMOS settings to a rescue disk for later recovery. In this chapter, you use the Internet to download a shareware utility to record CMOS settings and later recover them.

1. Access the Internet and then go to this address: *www.shareware.com*. Search on "CMOS" to list the various shareware utilities available. Select and download CMOS.ZIP. You can then exit the Internet.
2. Explode the compressed file and print the CMOS.TXT documentation file. Three utility programs are included:
   - CMOSSAVE.COM saves the CMOS settings to a file.
   - CMOSCHK.COM compares the CMOS settings to the last saved version.
   - CMOSREST.COM restores the CMOS settings from the file.
3. Access a DOS prompt and save the CMOS settings to a file on a floppy disk, using this command:

   ```
   CMOSSAVE.COM A:\MYFILE.SAV
   ```
4. Compare the settings stored in the file to the current CMOS settings, using this command:

   ```
   CMOSCHK.COM A:\MYFILE.SAV
   ```

   The results of these commands are shown in Figure 3-36.

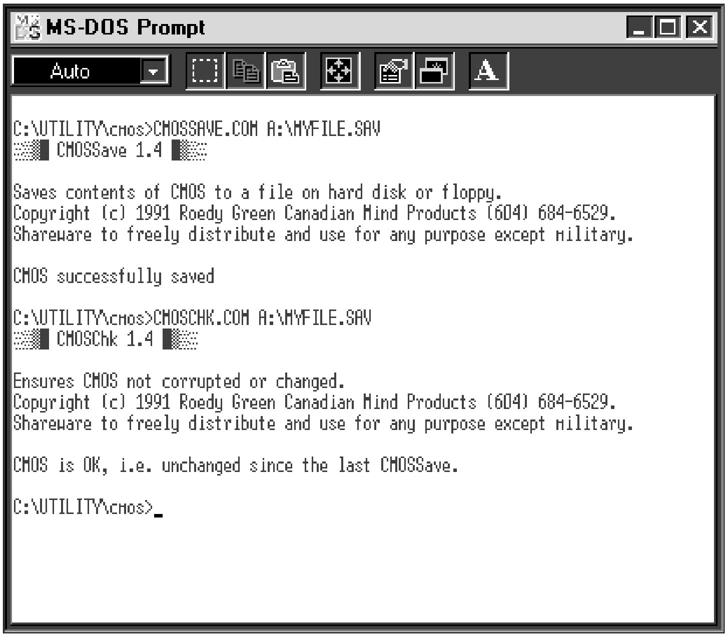

Figure 3-36 Using a shareware utility to save CMOS settings

Hands-on Project

Using a System Board Diagnostic Utility

A well-known diagnostic utility to help solve computer problems is AMIDiag from American Megatrends, Inc. The utility is DOS-based and works under both DOS and Windows 9x. You can download the utility from the Internet.

1. Access the Internet and then go to this address: *www.shareware.com*.
2. Locate the Quick Search text box and search on "amidiag" (don't enter the quotation marks). Download amidiag.zip to your PC. This file is a shareware version of AMIDiag for PC diagnostics.
3. Leave the Internet and expand the file by double-clicking it in either Windows 3.x or Windows 9x.
4. At a DOS prompt, change to the directory where the demo software files are stored.
5. Execute the first program by entering this command: AMIDIAG
6. The screen shown in Figure 3-37 should appear. Perform the test of processor speed. What is the detected speed?

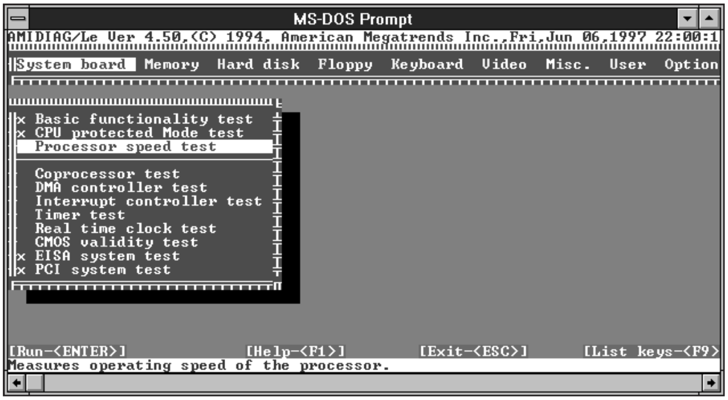

Figure 3-37 AMIDiag opening menu

7. On the Memory menu, perform all the tests that this demonstration version of the software allows. Record any errors detected.
8. On the Misc menu, perform the serial port test. Write down any error messages that you get. If you get an unexpected error, perform the test more than once. Do you get the same results each time?
9. On the Options menu, select System Information. If you received errors in the test above, this program might lock up, and you might need to reboot. If you complete the information check successfully, write down the results.
10. On the system board menu, select DMA Controller Test. Why doesn't this test work?
11. Exit the program, returning to the DOS prompt.

Practice Activity

Using old or defective expansion cards and system boards, practice inserting and removing expansion cards and chips.

Print a Summary of Your System Hardware

1. In Windows 9x, right-click the **My Computer** icon.
2. On the shortcut menu, select **Properties**.
3. Click the **Device Manager** tab.
4. View devices by type.
5. Click the **Ports** (Com & LPT) icon.
6. Click the **Print** button.
7. Print **Selected class or device**.

Understand Hardware Documentation

Obtain the manual for the system board for your PC. (If you can't find the manual, try downloading it from the system-board manufacturer's web site.) List at least three functions of jumpers on the board as well as the corresponding jumper numbers.

Learn to Use the Windows 9x Help Feature

1. In Windows 9x, click the **Start** button on the Taskbar.
2. Click **Help**.
3. Click the **Contents** tab.
4. Double-click **Troubleshooting**.
5. Select and print **If you have a hardware conflict**.

Troubleshoot Setup Errors

Have someone change the DIP switches, jumpers, or CMOS configuration on a system board. (First make sure someone records the original settings, and don't attempt this project unless you have the system-board documentation.) Troubleshoot the system. Use this opportunity to learn to take notes as you work. List each error you encounter and what you did to work toward a solution.

Use the Internet for Research

1. Search the web site of Intel, AMD, or Cyrix and print information on the most recent offering of a CPU for microcomputers offered by the company.
2. Using your own or a lab computer, pretend that the system-board manual is not available, and you need to know the settings for the jumpers on the system board. Identify the manufacturer of the system board and research the web site for that manufacturer. Print the jumper settings for the system board available on the web site.
3. Does your current system BIOS support USB? To find out, go to the USB web site at *www.usb.org*. Have the BIOS manufacturer and version available.
4. Prepare a presentation about CPUs. Include the history of microprocessors, how a CPU is made, and how a CPU works. Use these web sites as resources:

   www.intel.com/education/mpuworks/index.htm

   www.intel.com/education/chips/index.htm

   www.intel.com/intel/museum/25anniv/index.htm

Research the Market

1. In a current computer magazine, find the speed and price of the fastest PC CPU on the market today.
2. In a current computer magazine, find the speed and price of the fastest PC RAM module on the market today.

Observe Hardware Conflict Errors

Have someone set up a troubleshooting practice problem by forcing two hardware devices on a PC to use the same IRQ. Troubleshoot the problem. Take notes as you go. Describe the errors that you see and what you do to solve the problem.